arXiv:1410.3234v1 [stat.ML] 13 Oct 2014

Markov Random Fields and Mass Spectra Discrimination

AO KONG

School of Finance, Nanjing University of Finance and Economics, Nanjing, China

ROBERT AZENCOTT*

Dept of Mathematics, University of Houston, TX, USA

Emeritus Prof, C.M.L.A. Ecole Normale Superieure, Cachan, France

razencot@math.uh.edu

Summary

For mass spectra acquired from cancer patients by MALDI or SELDI techniques, automated discrimination between cancer types or stages has often been implemented by machine learnings. These techniques typically generate "black-box" classifiers, which are difficult to interpret biologically. We develop new and efficient signature discovery algorithms leading to interpretable signatures combining the discriminating power of explicitly selected small groups of biomarkers, identified by their m/z ratios. Our approach is based on rigorous stochastic modeling of "homogeneous" datasets of mass spectra by a versatile class of parameterized Markov Random Fields. We present detailed algorithms validated by precise theoretical results. We also outline the successful tests of our approach to generate efficient explicit signatures for six benchmark discrimination tasks, based on mass spectra acquired from colorectal cancer patients, as well as from ovarian cancer patients.

Key words: MALDI/SELDI data; ovarian cancer; colorectal cancer; biomarker signature discovery; Markov Random Fields; autologistic distribution.

*To whom correspondence should be addressed.

1. Introduction

In proteomics, mass spectrometry is a broadly used protein profiling technology to study the mixture of proteins/peptides present in biological tissues or fluids, and is an efficient tool for identification of cancer type and stage (Eckel-Passow *and others* (2009)).

Mass spectrometry can involve two soft ionization techniques: matrix-assisted laser desorption ionization (MALDI) and surface-enhanced laser desorption and ionization (SELDI). For each analyzed fluid sample, MALDI or SELDI hardwares generate a high-dimensional mass spectrum, recording between 10,000 and 20,000 "mass-to-charge (m/z) ratios" corresponding to the ionized peptides present in the fluid sample, as well as "intensities" roughly quantifying the concentrations of these peptides in the sample. Generally m/z ratios take values anywhere between 200 and 20,000 Daltons, and are acquired with a known *relative* accuracy ρ which depends on the acquisition modalities, and ranges from 0.1% to 0.3%.

Analyzing this type of high dimensional data oftern requires specialized software tools, implementing sophisticated machine learning techniques such as SVM (support vector machines) (Li *and others* (2004), Yu *and others* (2005)), artificial neural networks (Ball *and others* (2002)), or random forests (Izmirlian (2004)). These techniques typically generate "black-box" classifiers, which often reach good discrimination levels between cancerous and control groups, but are difficult to interpret biologically in terms of characteristic biomarkers patterns. This often leads to unexpected performance variations on totally new data sets. To develop clinically usable software tools for analysis of mass spectra acuired by MALDI or SELDI hardwares, a key step is to implement automated discovery of explicit "signatures", i.e. short lists of proteomic biomarkers with high discriminating powers between cancer groups (Yasui *and others* (2003)). Some easily interpretable automatic classifiers, such as linear combinations of biomarker weights (Wang and Chang (2011)), can be found in previous studies, but these approaches do not attempt to quantify the discriminating impact of simultaneous presence for specific pairs of biomarkers.

In this paper, we generate easily interpretable biomarker signatures discriminating between two arbitrary but homogeneous groups of mass spectra G^+ and G^- by stochastic modeling of biomarkers interactions, taking precisely into account the co-activity of pairs of biomarkers. To this end, we fit parametric Markov Random Fields (MRFs) π^+ and π^- to G^+ and G^- by Maximum Pseudo-Likelihood Estimation (MPLE). Recall that MRFs have been successfully used to

model spatial dependencies in high dimensional interacting systems as well as in image and signal analysis (Chalmond (2003), Brémaud (1999), Azencott and Graffigne (1992)). Discrimination between G^+ and G^- is then achieved by computing the optimal separator between the probability distributions π^+ and π^-.

We have studied quite precisely the asymptotic performance of our approach for large data samples, and we have successfully benchmarked our MRF based signature discovery technique on MALDI and SELDI data sets respectively acquired from colorectal and ovarian cancer patients.

2. Benchmark Mass Spectra Datasets

2.1 *Ovarian Cancer Data "4-3-02" and "8-07-02":*

These two mass spectra data sets, acquired by SELDI-TOF techniques have been previously preprocessed and studied by other authors in Assareh and Moradi (2007), Zhu *and others* (2003), Alexe *and others* (2004) and can be freely downloaded from the NCI-FDA clinical proteomics databank (http://home.ccr.cancer.gov/ncifdaproteomics/ ppatterns.asp).

The "4-3-02" set includes a Control group (CTR) of 116 mass spectra and an Ovarian Cancer (OVC) group of 100 mass spectra. The "8-07-02" set includes a Control group (CTR) of 91 mass spectra and an Ovarian Cancer (OVC) group of 162 mass spectra. As available online, these mass spectra are already "aligned" to 15,154 reference m/z ratios ranging from 0 to 20,000 Daltons, and the relative accuracy of m/z ratios is $\rho = 0.1\%$.

2.2 *Colorectal Cancer Data:*

For this newly published data set in our previous study Kong *and others* (2014), plasma samples from colorectal cancer patients and a control group were provided by 1st Surgical Clinic, Dept of Surgical, Oncological and Gastroenterological Sciences at University of Padova (Italy). The mass spectra were then acquired through MALDI-TOF techniques by A. Bouamrani, E. Tasciotti, M. Ferrari (Dept. of Nanomedicine, The Methodist Hospital Research Institute, Houston, USA).

This set includes an Adenoma group (ADE) of 54 spectra, an Early Colorectal Cancer group (ECR) of 80 spectra, a Late Colorectal Cancer group (LCR) of 74 spectra and a Control group (CTR) of 30 spectra. The union of ADE, ECR, LCR will be called the Colorectal Cancer (CRC)

Group. All these mass spectra were generated by MALDI-TOF technique, with a relative accuracy of $\rho = 0.3\%$ on m/z ratios, which range from 800 to 10,000 Daltons.

2.3 *Discrimination Tasks*

Our paper presents new signature discovery algorithms based on Markov Random Fields modeling of mass spectra. We have then applied and evaluated our MRF discrimination techniques to the differentiation between cancer stages, as well as between cancer and control. So we have tested implementation and performances of our MRF approach on 6 benchmark discrimination tasks, using the three sets just described of mass spectra acquired from cancer patients :
1) ADE vs ECR, ADE vs LCR, ECR vs LCR, CRC vs CTR for the colorectal cancer dataset,
2) OVC vs CTR for each one of the two ovarian cancer datasets.

3. Binary Coding of Mass Spectra

3.1 *Pre-processing of Proteomic Mass Spectra*

Pre-processing is an important procedure to lower mass spectra dimensionality and to remove acquisition noise, which affects both m/z ratios and intensities. When an intensity peak is detected at abscissa A, its true m/z ratio could lie anywhere within $[\, A(1-\rho), A(1+\rho)\,]$, where the relative accuracy ρ is determined by the acquisition hardware.

For better context control, we apply to each raw mass spectrum our own sequence of pipelined classical pre-processing steps: normalization, smoothing, noise extraction, baseline removal, peak detection as outlined in our previous paper (Kong *and others* (2014)). On each mass spectrum, these pre-processing steps detect a usually long list of "strong intensity peaks". The m/z abscissas of these detected peaks indicate peptides which could potentially be biomarkers strongly discriminating between cancer types or stages .

To condense all these approximate peak abscissas, we generate a fixed list of "reference biomarkers" B_s with m/z abscissas $B_s = B_1(1+\rho)^{s-1}, 1 \leqslant s \leqslant L$, where B_1 and B_L are the smallest and the largest m/z ratios among all spectra in our dataset. Then $S = \{\, 1, \ldots, L\,\}$ will be called our "set of sites", and $\mathcal{B} = \mathcal{B}(L)$ will be the set of all binary vectors x of length L, with coordinates x_s indexed by S.

3.2 *Binary Coding of Mass Spectra*

We say that a reference biomarker B_s is "activated" by a mass spectrum M if and only if at least one detected peak of M is positioned within the window of $[\ B_s - \rho B_s,\ B_s + \rho B_s\ ]$. Each mass spectrum M can then be "coded" by a binary vector $x = x(M) \in \mathcal{B}(L)$ as in Kong *and others* (2014): for each site $s \in S$, we set $x_s(M) = 1$ if B_s is activated by M and $x_s(M) = 0$ otherwise.

Any group G of n mass spectra thus generates the set $BinG$ of n binary vectors $x(M) \in \mathcal{B}(L)$ coding as just indicated the mass spectra M of G.

4. Binary Markov Random Fields and Autologistic Distributions

We now want to consider the data set $BinG$ as a sample of n independent observations of a *random* binary vector X taking values in the set of binary vectors $\mathcal{B} = \mathcal{B}(L)$.

We will systematically model the unknown probability distribution P of X by a Markov Random Field on $\mathcal{B}$. A first statistical analysis of the data set $BinG$ outlined in Section 8.2 below will identify for each site $s \in S$ the subset $\mathcal{N}_s$ of all sites $t \in S - s$ for which we "expect" the coordinate pairs (X_s, X_t) to be strongly correlated.

Recall that a binary random vector $X = \{X_s\}_{s\in S}$ with values in $\mathcal{B}$ is called a Markov Random Field (MRF) with respect to the family $\mathcal{N}_s$ if for all sites $s \in S$

$$P\{X_s \mid X_{S\setminus s}\} = P\{X_s \mid X_{\mathcal{N}_s}\},$$

where for any $K \subset S$, we denote by X_K the set of random variables $\{X_s\}_{s\in K}$. The distribution P of X then belongs to the family of Gibbs distributions, and can be described concretely through its system of "cliques". Recall that a clique is a subset C of S, such that $t \in \mathcal{N}_s$ for all distinct pairs $\{s, t\} \in C$. All single sites are then cliques of cardinal 1.

Here, we focus on *autologistic* distributions, which are the Gibbs distributions for which all cliques have cardinal $\leqslant 2$. They are naturally parameterized by a vector space Θ isomorphic to R^k, with $k = L(1 + L)/2$, and where the coordinates of any parameter vector $\theta \in \Theta$ are denoted by θ_s and $\theta_{s,t}$ with $s, t \in S$ and $s < t$.

For each $x \in \mathcal{B}$, let $U(x) \in \Theta$ be the vector with coordinates

$$U_s(x) = x_s\ , \quad U_{s,t}(x) = x_s x_t\ , \quad \textit{for} \quad s, t \in S,\ s < t.$$

The scalar product of θ and $U(x)$ in Θ is then

$$< \theta, U(x) >= \sum_{s \in S} \theta_s x_s + \sum_{s<t} \theta_{s,t} x_s x_t.$$

For each $\theta \in \Theta$, the autologistic distribution π_θ is defined for all $x \in \mathcal{B}$, by

$$\pi_\theta(x) = \frac{1}{Z(\theta)} e^{-<\theta, U(x)>}, \tag{4.1}$$

where $Z(\theta) = \sum_{y \in \mathcal{B}} e^{-<\theta, U(y)>}$ is the partition function.

5. Fitting Autologistic Distributions to Mass Spectra data sets: Maximum Pseudo-Likelihood Estimators (MPLE)

As seen in Section 3.2, the binary coding of a mass spectra data set G generates a set $BinG$ of n binary vectors of length L, and in concrete applications to cancer data, L is typically much larger that n. Reliable fitting of an autologistic distribution π_θ to $BinG$ then requires a strong dimension reduction from L sites to a much smaller set $S(d)$ of d sites adequatedly selected in S. This is achieved by the "Feature Selection" algorithm we present further on in Section 8.1. Restricting each binary vector $x \in BinG$ to the d sites in $S(d)$ transforms $BinG$ into a set BG of n binary vectors of length d.

Defining an autologistic distribution on $\mathcal{B}(d)$ involves selecting a specific family of pairs of sites $\{s,t\}$ with $s,t \in S(d)$, for which the binary random variables X_s, X_t are expected to have sizeable correlation. The maximum number $d(d-1)/2$ of these potential cliques of order 2 is often still too high with respect to n. So we seek model robustness by enforcing parameter parsimony, which leads us to retain only a moderate number $c << d(d-1)/2$ of pertinent cliques of order 2. This is implemented by the "Clique Discovery" procedure outlined in Section 8.2.

After selecting d sites and a set $\mathcal{C}$ of c cliques of order 2 in $S(d)$, we seek to model the data set $BG \subset \mathcal{B}(d)$ of n binary vectors by an autologistic probability distribution π_θ of the form (4.1), where we now impose on θ the constraints $\theta_s = 0$ whenever s is not in $S(d)$ and $\theta_{s,t} = 0$ whenever $\{s,t\}$ is not in $\mathcal{C}$, so that the unknown parameter vector θ is now forced to belong to a precise vector subspace of Θ, of dimension $(c+d)$.

To achieve this model fitting to data, θ must be estimated from the n data. After comparative testing of several classical estimation techniques on our benchmark examples, we have implemented all our model fitting to data through Maximum Pseudo-Likelihood Estimators (MPLEs).

These estimators were introduced by Besag (1975) and have played an important role in parameter estimation of spatial models before the current intensive use of Monte Carlo methods. Indeed computing the MPLE requires no simulation of random Gibbs configurations, leading to fast computing speed.

5.1 *Pseudo-Likelihood*

As just seen, the preceding selection of d sites s and c cliques $\{s,t\}$ of order 2 forces a precise set of coordinates of $\theta \in \Theta$ to be equal to 0. The estimation of the non zero coordinates of θ will be based on maximizing the average pseudo-likelihood of the observed data. We now recall how one computes pseudo-likelihoods. For brevity and to simplify notations, we deliberately restrict our theoretical presentation to the case where no constraints are imposed on the coordinates $\theta_s, \theta_{s,t}$ of θ. The constrained case is an easy extension of the non-constrained case.

Let x be any observed binary vector and let $T \in \Theta$ be any tentative estimate of θ. The *pseudo-likelihood* $PL(x,T)$ is classically defined as the product of all "local specifications" under the tentative autologistic distribution π_T

$$PL(x,T) = \prod_s \pi_T(x_s|x_{S-s}) = \prod_s P_T(Y_s = x_s \mid Y_{S-s} = x_{S-s}),$$

where Y is a random binary configuration with distribution π_T. Consider the linear functions of $T \in \Theta$ defined by

$$a_{s,x}(T) = -\,(T_s + \sum_{s<r} T_{sr}\, x_r) = < A(s,x), T >, \tag{5.2}$$

where the vector $A(s,x) \in \Theta$ has coordinates

$$A_s(s,x) = -1, \quad A_{sr}(s,x) = -x_r, \quad \textit{for all} \quad s < r. \tag{5.3}$$

Define two functions of z in R by

$$g(z) = 1/(1+e^z), \quad h(z) = e^z/(1+e^z). \tag{5.4}$$

The conditional specification of Y_s under π_T can then be written

$$p_{s,x}(T) = P_T(Y_s = x_s \mid Y_{S-s} = x_{S-s}) = x_s\, h(a_{s,x}(T)) + (1-x_s)\, g(a_{s,x}(T)).$$

Since x_s is either 0 or 1, this easily implies

$$LPL_{s,x}(T) = \log(p_{s,x}(T)) = x_s\, a_{s,x}(T) + \log(\, g(a_{s,x}(T))\,), \tag{5.5}$$

the log pseudo-likelihood function LPL is hence given by

$$LPL(x,T) = \sum_{s\in S} LPL_{s,x}(T) = \sum_{s\in S} \left[\, x_s\, a_{s,x}(T) + \log(g(a_{s,x}(T))\,\right].$$

5.2 *Computation of the MPLE*

Let X be a random configuration with autologistic distribution π_θ. For all $T \in \Theta$, define the *mean log pseudo-likelihood* $g_\theta(T)$ by

$$g_\theta(T) = E_\theta(LPL(X,T)) = \sum_{s\in S}\sum_{x\in\mathcal{B}} \pi_\theta(x)\, \left[\, x_s\, a_{s,x}(T) + \log(g(a_{s,x}(T))\,\right]. \tag{5.6}$$

The theoretical principle of the MPLE algorithm is to seek a vector estimate T of θ which maximizes in T the *mean log pseudo-likelihood*. This approach relies on the strict concavity of $g_\theta(T)$ as a function of T (see Proposition 1 below).

Consider n independent observed configurations $\{X^1,\ldots,X^n\}$, generated by the same unknown autologistic distribution π_θ. Due to the law of large numbers one can approximate the unknown function $g_\theta(T) = E_\theta(LPL(X,T))$ by the *empirical log pseudo-loglikelihood*

$$\hat{g}(T,n) = \frac{1}{n}\sum_{j=1}^{n} LPL(X^j,T). \tag{5.7}$$

Let $\partial_T LPL(x,T)$ and $Hess_T LPL(x,T)$ denote the gradient and Hessian matrices of $LPL(x,T)$ with respect to T, then the gradient and Hessian matrix of $\hat{g}(T,n)$ with respect to T are

$$\hat{G}(T,n) = \frac{1}{n}\sum_{j=1}^{n} \partial_T LPL(X^j,T), \tag{5.8}$$

and

$$\hat{H}(T,n) = \frac{1}{n}\sum_{j=1}^{n} Hess_T LPL(X^j,T). \tag{5.9}$$

The Maximum Pseudo-Likelihood Estimator $\hat{\theta}(n)$ of θ is defined as a vector which minimizes $\hat{g}(T,n)$ in T, and hence verifies the non linear vector equation

$$\hat{G}(\hat{\theta}(n),n) = 0.$$

The existence and fast computability of $\hat{\theta}(n)$ is due to the following proposition.

Proposition 1 For each $\theta \in \Theta$, the mean log pseudo-likelihood $g_\theta(T)$ is a strictly concave function of $T \in \Theta$, and reaches its maximum in T at the unique point $T = \theta$. Moreover, the empirical pseudo-likelihood $\hat{g}(T, n)$ is also concave in $T \in \Theta$, and becomes almost surely strictly concave as $n \to \infty$.

Proof: The proof is given in Appendix A of the Supplementary Materials.

Due to the concavity of $\hat{g}(T, n)$, we implemented a standard gradient descent to generate a sequence $T(j)$ converging to $\hat{\theta}(n)$ as j increases

$$T(j+1) = T(j) + \epsilon \, \hat{G}(\, T(j), n \,).$$

One stops iterating when $|| \, \hat{G}(\, T(j), n \,) \, ||$ becomes inferior to a user-chosen small threshold.

For each one of our benchmark discrimination studies, autologistic modeling by MPLE was used intensively to parametrize roughly around 2,000 to 5,000 Gibbs models of dimension less than 30, in order to explore enough potentially discriminating combinations of sites and cliques of order 2, automatically selected among the large list of reference biomarkers. Each MPLE modeling was implemented by gradient descent involving 200 iteration steps, with a step size $\epsilon = 0.05$. On a 1.3 GHz MacOS PC, the average computing time per Gibbs model was about 1 second for colorectal cancer stage groups of mass spectra, which all had small size $\leqslant 80$; this CPU time increased to about 8 seconds for groups of size ~ 100, such as our ovarian cancer groups, and reached about 15 seconds for groups of size ~ 200, such as the full colorectal cancer group.

5.3 *Asymptotic Normality of MPLE*

Most early results on asymptotic normality of the MPLE have focused on Gibbs random fields on infinite lattices under Dobrushin unicity conditions (see for instance Jensen and Kunsch (1994)). For Gibbs random fields on the finite set $\mathcal{B}$ of binary vectors of fixed length L, the Dobrushin conditions become irrelevant, so that many publications consider asymptotic consistency and normality of MPLEs as valid for finite configurations spaces, but without refering to explicit proofs. We state a precise asymptotic result, proved in the mathematical Appendix B of the Supplementary Materials.

Proposition 2 For any autologistic distribution π_θ on the set of binary configurations $\mathcal{B}$, the MPLE

estimators $\hat{\theta}(n)$ of θ are asymptotically consistent as the number of observations $n \to \infty$. The normalized vectors of estimation errors $\sqrt{n}\ (\hat{\theta}_n - \theta)$ are asymptotically Gaussian with mean zero and a covariance matrix $\Gamma(\theta)$ computable as indicated in the Appendix B of the Supplementary Materials.

Proof: The proof is given in Appendix B of the Supplementary Materials.

5.4 *Elimination of non-significant parameters*

The preceding asymptotic normality result provides a tool to decide if some estimated parameter coordinates should be replaced by 0. Indeed, the diagonal elements Γ_s and $\Gamma_{s,t}$ of the covariance matrix $\Gamma(\theta)$ can easily be approximated from the data (see Appendix B). The standard deviations $\hat{\theta}_s(n) - \theta_s$ and $\hat{\theta}_{s,t}(n) - \theta_{s,t}$ are then $\sigma_s(n) \sim \sqrt{\Gamma_s}/\sqrt{n}$ and $\sigma_{s,t}(n) \sim \sqrt{\Gamma_{s,t}}/\sqrt{n}$, which yields explicit Gaussian 90% confidence intervals for θ_s and $\theta_{s,t}$. Whenever one of these 90% confidence intervals contains 0, we constrain the corresponding coordinate of $\hat{\theta}$ to be 0. The autologistic model is then re-estimated by MPLE, but the gradient descent implementing MPLE computation now takes into account the complementary constraints just introduced on θ. This procedure is iterated until all estimated parameters are significantly different from 0.

In our benchmark studies, the sizes $n = 54, 80, 74, 30$ of the four colorectal cancer data sets ADE, ECR, LCR CTR were rather small. So in the autologistic modeling of these 4 cancer datasets the relative accuracies $\sqrt{\Gamma_s}/(\sqrt{n}\hat{\theta}_s)$ of several estimated coordinates $\hat{\theta}_s$ ranged between 30% and 50% (see for instance the accuracy results displayed in Table 4 below). It would clearly have been desirable to double the sizes of our colorectal cancer data sets, but the mass spectra acquisition phase had already been fully terminated when the present statistical study began. Our four ovarian cancer groups have larger sample sizes $n = 116, 100, 91, 162$ which naturally led to more accurate modeling. Each one of our benchmark studies involved the estimation of at least two thousand autologistic models, and a quick analysis of all the associated estimation accuracy results showed that "ideal" datasets sizes $n \sim 300$ would be amply sufficient for very accurate modeling.

Consider two groups G^+ and G^- of binarized mass spectra, with resp. sizes n^+ and n^-. For automatic discrimination between these two groups, we will model both of them by autologistic

distributions π^+ and π^-. Modeling accuracy will not be the main criterion , since the goal is then to identify *short* biomarker signatures enabling high performance discrimination. The well known parameter parsimony principle is here quite relevant, and suggests to strongly restrict the total number m of nonzero coordinates in the joint parametrization of π^+ and π^-. Ideally one should have $m << n^+ + n^-$. In our benchmark discrimination tasks, our optimized signature selection procedure typically yielded $5 \leqslant m \leqslant 30$ and $m \leqslant (n^+ + n^-)/5$.

5.5 *Quality of Fit: Empirical Estimation*

After fitting an autologistic distribution $\nu = \pi_{\hat{\theta}}$ to a data set G of n observed binary vectors, the *quality of fit* of ν to the data can be evaluated as follows.

First compute the empirical log-likelihood LL of the data set G under the probability ν. To approximate the distribution λ of LL under ν, we implement a classical Gibbs sampler (Brémaud (1999)) based on the Gibbs distribution ν to simulate 1000 virtual data sets $\mathcal{G}_1, \mathcal{G}_2, \ldots, \mathcal{G}_{1000}$, where each $\mathcal{G}_j$ contains n random binary vectors having the same probability distribution ν. Then compute the log likelihood LL_j of each virtual data set $\mathcal{G}_j$ under the probability ν. After re-ordering the list of 1000 log-likelihood values $LL_1, LL_2, \ldots, LL_{1000}$, the rank of LL becomes $1 \leqslant r(LL) \leqslant 1000$.

For the log likelihood distribution λ the number $Q = r(LL)/1000$ is the (random) quantile corresponding to the observed LL. Under the true but unknown probability π_θ and for n large, the quantity $\sqrt{n}\,(Q - 50\%)$ is approximately Gaussian with mean zero, and hence can easily be used to quantify the goodness of fit of $\pi_{\hat{\theta}}$ to the data set G. Namely, a good quality of fit between $\pi_{\hat{\theta}}$ and G should correspond to a Q value close to 50%.

In our benchmark studies, the autologistic distributions estimated by MPLE generally have very good quality of fit to real mass spectra datasets. For example, in our colorectal cancer data, to discriminate the patients groups $G = LCR$ from ECR, the algorithms described below identified the "best" autologistic model ν for the group G (which contained $n = 74$ mass spectra). The model ν involved only 18 optimally selected biomarkers and the quantile $Q = r(LL)/1000$ had the value 51.3%, which indicates a very good quality of fit. Figure 1 displays the histogram of 1000 virtual log likelihood values LL_j generated by Gibbs sampler based on ν, the bold vertical black line has abscissa $Q = 51.3\%$.

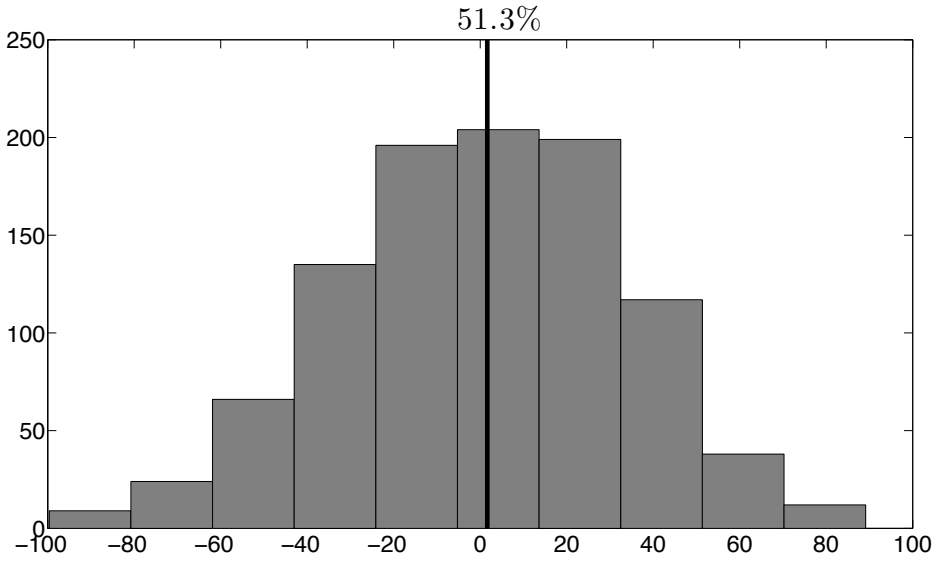

Fig. 1. This figure displays the quality of fit of an autologistic model $\nu = \pi_{\hat{\theta}}$ fitted to a data set of 74 mass spectra, namely the colorectal cancer group LCR. The 18 biomarkers involved in this model were selected (see details further on) to optimize the discrimination ECR vs LCR. The log likelihood histogram displayed here is based on 1000 virtual log likelihood values resp. computed on 1000 simulated random samples of 80 binary vectors of length 18, simulated from ν by Gibbs sampler. The bold vertical black line correspond to the true log likelihood value computed on the actual 74 binarized spectra in LCR.

6. Optimal Discrimination between Autologistic Models

On a set $\mathcal{B}$ of binary vectors with length L, let $\pi^+ = \pi_{\theta^+}$ and $\pi^- = \pi_{\theta^-}$ be two autologistic distributions, parameterized by $\theta^+, \theta^- \in \Theta$, and with partition functions $Z^+ = Z(\theta^+)$ and $Z^- = Z(\theta^-)$. We are using the same set of cliques for π^+ and π^-(namely, π^+ and π^- share the same $U(x)$), which is not a restriction since the coordinates of θ^+ and θ^- are allowed to take the value zero.

Call *decision rule* any real valued function g defined on $\mathcal{B}$. Each such g classifies random observation $x \in \mathcal{B}$ as "generated by π^+" if $g(x) > 0$, and as "generated by π^-" if $g(x) < 0$. The *performance* of any decision rule g is quantified by the two probabilities $p^+(g)$ and $p^-(g)$ of correct decisions when g is confronted to random configurations respectively generated either only by π^+ or only by π^-, so that

$$p^+(g) = \pi^+[\, x \mid g(x) > 0\,] \quad \textit{and} \quad p^-(g) = \pi^-[\, x \mid g(x) < 0\,].$$

Optimal decision rules are characterized as follows.

Optimal Decision Rules: Fix any weight coefficient $0 < \alpha < 1$. Among all decision rules g, there is an *optimal decision rule* f which maximizes $(\alpha\, p^+(g) + (1-\alpha)\, p^-(g))$. More precisely, there is then a number γ determined by α, such that for all $x \in \mathcal{B}$, the optimal f verifies

$$f(x) = < u, U(x) > + a, \quad \textit{with} \quad u = \theta^- - \theta^+ \quad \textit{and} \quad a = \log(Z^-/Z^+) - \gamma.$$

When $\alpha = 1/2$, the *performance criterion* to be maximized becomes $PERF(g) = (p^+(g) + p^-(g))/2$ and the number γ is actually zero.

Proof of this characterization: The relation $f(x) > 0$ is equivalent to $D(x) > e^\gamma$ where the likelihood function $D(x)$ is the density function $\frac{d\pi^+}{d\pi^-}(x)$. The announced result can then be derived from the Neyman-Pearson lemma, essentially as was done in Proposition 5.2 of Azencott *and others* (2014). When $\alpha = 1/2$, the symmetry between π^+ and π^- immediately shows that $\gamma = 0$.

7. Discrimination errors due to estimation errors on model parameters

7.1 *Estimated optimal separator*

In concrete discrimination tasks, the autologistic models π^+ and π^- are not known but derived from independent observed configurations $X^1, \ldots, X^n$ and $Y^1, \ldots, Y^n$, which we consider as separately generated by π^+ and π^-. The unknown vectors of parameters θ^+ and θ^- are then naturally replaced by their MPLE estimators $\hat\theta_n^+$ and $\hat\theta_n^-$. In our benchmark applications below, we systematically maximize the simplest *performance criterion PERF(f)* $= (p^+(f) + p^-(f))/2$, so that due to the "Optimal Decision Rules" in Section 6, the optimal separator f is then given by

$$f(x) = < u, U(x) > + a, \quad \textit{with} \quad u = \theta^- - \theta^+ \quad \textit{and} \quad a = \log(Z^-/Z^+). \tag{7.10}$$

After replacing the parameters of (π^+, π^-) by their MPLE estimates, the unknown models and their partition functions are replaced by their estimates $(\hat\pi^+, \hat\pi^-)$ and $(\hat Z^+, \hat Z^-)$. The optimal separator f is estimated by $\hat f$ where

$$\hat f(x) = < \hat u, U(x) > + \hat a \quad \textit{with} \quad \hat u = \hat\theta_n^- - \theta_n^+, \quad \textit{and} \quad \hat a = \log(\hat Z^-/\hat Z^+). \tag{7.11}$$

When the classification of observed binary vectors x is based on the sign of $\hat f(x)$ instead of

sign($f(x)$), the performance quantifiers $p^+(f)$ and $p^-(f)$ are replaced by $\hat{p}^+$ and $\hat{p}^-$ given by

$$\hat{p}^+ = \pi^+ [\, x \mid \hat{f}(x) > 0 \,] \quad \textit{and} \quad \hat{p}^- = \pi^- [\, x \mid \hat{f}(x) < 0 \,].$$

We are now going to estimate the errors $(\hat{p}^+ - p^+)$ and $(\hat{p}^- - p^-)$.

Proposition 3 In the preceding situation, as $n \to \infty$, the normalized error vector $\sqrt{n}(\hat{u} - u)$ is asymptotically Gaussian with mean zero and covariance matrix $Cov = \Gamma(\theta^+) + \Gamma(\theta^-)$, where the matrices $\Gamma(\theta)$ have been computed in Proposition 2. Moreover, $\sqrt{n}\,(\hat{a} - a)$ is also asymptotically Gaussian with mean zero and asymptotic variance $var(\theta^+, \theta^-)$ verifying

$$\mathrm{var}(\theta^+, \theta^-) \leqslant \frac{\mathrm{L}(1+\mathrm{L})}{2} \times \mathrm{trace}\left[\Gamma(\theta^+) + \Gamma(\theta^-)\right]. \tag{7.12}$$

Proof: The proof is given in Appendix C of the Supplementary Materials.

We now compare the discriminating powers of decision regions defined by two affine separators.

Proposition 4 On the set $\mathcal{B}$ of binary vectors x with length L, consider the two separators

$$f(x) =< u, U(x) > + a\ , \quad \phi(x) =< \eta, U(x) > + \alpha$$

parametrized by $u, \eta \in \Theta$ and $a, \alpha \in R$. Fix any autologistic distribution π_θ on $\mathcal{B}$. Let $p(f)$ and $p(\phi)$ be the two probabilities

$$p(f) = \pi_\theta[\, x \mid f(x) > 0 \,] \quad \textit{and} \quad p(\phi) = \pi_\theta[\, x \mid \phi(x) > 0 \,].$$

Recall that $k = L(1+L)/2$. Fix u, a, θ and assume that $0 < p(f) < 1$. For any arbitrary small percentage $0\% \leqslant \gamma < 100\%$, there is a strictly positive number q depending only on θ, u, a, γ such that for any pair (η, α) verifying

$$||\eta - u|| < \frac{q}{2\sqrt{k}} \quad \textit{and} \quad |\,\alpha - a\,| < q/2, \tag{7.13}$$

one must then have $|\,p(f) - p(\phi)\,| \leqslant \gamma$.

Proof: The proof is given in Appendix C of the Supplementary Materials.

7.2 *Control of discrimination errors due to separator estimation*

Proposition 5 Fix two autologistic distributions $\pi^+$and π^- parameterized by $\theta^+, \theta^- \in \Theta$. Let $f(x) =< u, U(x) > + a$ be the optimal separator between π^+ and π^-, given by equation (7.10).

From two random samples of n configurations resp. generated by π^+ and π^-, one computes by equation (7.11) the estimator $\hat{f}(x) =< \hat{u}, U(x) > + \hat{a}$ of $f(x)$.

Let $p^+(f), p^-(f)$ and $p^+(\hat{f}), p^-(\hat{f})$ be the probabilities of correct discrimination between π^+ and π^- resp. achievable by the separators f and $\hat{f}$. Assume that $0 < p^+(f) < 1$ and $0 < p^-(f) < 1$.

Let γ and κ be two arbitrary small numbers verifying $0 \leqslant \gamma < 1$ and $0 < \kappa < 1$. Then one can find N such that for $n > N$ one has

$$P(\ |\, p^+(\hat{f}) - p^+ \,| \leqslant \gamma \ \ and \ \ |\, p^-(\hat{f}) - p^- \,| \leqslant \gamma\ ) \ \geqslant\ 1 - \kappa \tag{7.14}$$

within the following proof.

Proof: The proof and a practical estimate of N are given in the Appendix D of the Supplementary Materials.

8. Markov Random Field (MRF) Discrimination

8.1 *Biomarker Selection to Discriminate Between two Groups of Mass Spectra*

Concrete datasets of mass spectra acquired from cancer patients usually involve several distinct patient groups, but are often of moderate sizes inferior to 100 spectra, as in our benchmark studies below. For practical discrimination between several patient groups, our pre-processing of raw mass spectra is done simultaneously for the data of all these groups. This generates a list of reference biomarkers B_s with $s \in S$, where S typically has large size L. For instance further below, in our colorectal cancer study, we have $L \sim 800$, and for the benchmark ovarian cancer data, pre-processing by other teams had yielded $L \sim 15,000$.

When fitting autologistic models π^+, π^- to two training sets $BinG^+$, $BinG^-$ of binarized mass spectra, one wants to achieve statistically *robust* fitting, as well as high discriminating power between π^+ and π^-. Among the set of reference biomarkers $B_s, s \in S$, one must hence select a subset of small cardinal d, and for each selected B_s, the presence or absence of B_s within any x in $BinG^+$ or $BinG^-$ should provide strongly discriminating information about the correct classification of x. This leads us to the following *feature selection* algorithm.

Given two training sets $BinG^+$ and $BinG^-$ of binarized spectra, both included in the set of

binary vectors $\mathcal{B}(L)$, we compute for each $s \in S$ the resp. frequencies $m^+(s)$ and $m^-(s)$ of the event $x_s = 1$ for $x \in BinG^+$ and for $x \in BinG^-$. To ensure a significant presence of each selected biomarker B_s in at least one of the two groups, we fix a minimal frequency threshold thr, taking account of the number n of training data, and we select in S the subset $\hat{S}$ of all sites s for which $\min\{m^+(s), m^-(s)\} \geqslant thr$. For our benchmark studies where $30 \leqslant n \leqslant 162$ was rather small, we systematically used $thr = 20\%$.

As in Kong *and others* (2014), we quantify the *Discriminating Power* of each site s in $\hat{S}$ by the ratio $DP(s) = m^+(s)/m^-(s)$. This empirical definition is easily justified when one discriminates between two autologistic distributions where all pairs of sites are weakly correlated. Highly discriminating biomarkers B_s tend to have either very high or very low $DP(s)$ values. Since we want to discover highly discriminating signatures involving only a small number d of biomarkers, we fix a moderate number H (approximately equal to 1% of cardinal($\hat{S}$) in all our benchmark applications), and we focus only on the H sites $s \in \hat{S}$ with the highest $DP(s)$ and the H sites with the lowest $DP(s)$.

Fix any two positive integers $d^+ \leqslant H$ and $d^- \leqslant H$, so that $d << L$. The choice of d^+, d^- will be optimized further on. Within the biomarkers B_s with $s \in \hat{S}$, we select as "G^+ biomarkers" the d^+ biomarkers B_s with the highest $DP(s)$, and as "G^- biomarkers" the d^- biomarkers B_s with the lowest $DP(s)$.

The union $S(d) \subset \hat{S}$ of the two sets of d^+ and d^- sites just selected has cardinal $d = d^+ + d^- << b$. Each binary vector x in $BinG^+$ or $BinG^-$ is then systematically *restricted* to the d sites of $S(d)$. This restriction generates two sets of binary vectors in $\mathcal{B}(d)$ which we will denote by BG^+ and BG^-.

8.2 *Clique Discovery for two Autologistic Models*

Given the data set BG^+ of n binary vectors, and any pair of sites s, t, one can easily compute the empirical joint frequencies of the four events $\{X_s = i, X_t = j\}$, where i and j have binary values 0 or 1. We use this 2×2 contingency table to quantify the stochastic dependency between X_s and X_t by a classical χ^2-statistic with one degree of freedom, denoted $\chi^2(s,t)$. At the 95% significance level, X_s and X_t are thus considered "dependent" iff $\chi^2(s,t) > 3.84$, and the pair $s, t.$ will then be retained as a potential clique.

The feature selection procedure used to generate BG^+ and BG^- (see Section 8.1) implies that whenever B_s is a G^- biomarker, the frequency $m^+(s)$ of the event $X_s = 1$ within the dataset BG^+ will typically be quite small. For all our benchmark studies, we had indeed $m^+(s) < 5\%$ for all G^- biomarkers B_s. Practically, for mass spectra acquired from cancer patients, datasets sizes n are moderate, of the order of 100. So when $m^+(s) < 5\%$, it is hopeless to obtain reliable estimates of the dependency statistics $\chi^2(s,t)$ for any t.

To achieve a reasonably robust fitting of an autologistic model π^+ to the dataset BG^+, the set $\mathcal{C}^+$ of cliques of order 2 for π^+ will hence be restricted to include only pairs $\{s,t\}$ such that $\chi^2(s,t) > 3.84$ and both B_s and B_t are G^+ biomarkers. Call $\mathcal{PTC}^+$ the set of all such potential cliques. The cardinal c^+ of $\mathcal{C}^+$ will successively be fixed at any value inferior to $card(\mathcal{PTC}^+)$, and will be optimized further on. Once we fix c^+, we retain in $\mathcal{C}^+$ precisely the c^+ cliques $\{s,t\}$ belonging to $\mathcal{PTC}^+$ and having the c^+ highest $\chi^2(s,t)$ values.

Similarly, let $\mathcal{PTC}^-$ be the set of pairs B_s, B_t of G^- biomarkers such that $\chi^2(s,t) > 3.84$, with $\chi^2(s,t)$ evaluated on BG^-. To fit π^- to BG^-, we tentatively fix any $c^- \leqslant card(\mathcal{PTC}^-)$, to be optimized later on, and the set of cliques $\mathcal{C}^-$ for π^- are the c^- cliques $\{s,t\}$ with highest $\chi^2(s,t)$ in $\mathcal{PTC}^-$.

8.3 *Numerical Autologistic Model Fitting*

Call $\mathcal{D}(H)$ the set of all quadruplets of integers $dim = (d^+, c^+; d^-, c^-)$ such that

$$1 \leqslant d^+ \leqslant H\,,\;\; 1 \leqslant d^- \leqslant H\,,\;\; 0 \leqslant c^+ \leqslant card(\mathcal{PTC}^+)\,,\;\; 0 \leqslant c^- \leqslant card(\mathcal{PTC}^-)\,, \tag{8.15}$$

where the moderate integer H was selected in Section 8.1. The set $\mathcal{D}(H)$ has size $\sim O(H^2)$ which increases quickly with H.

For each quadruplet $dim = (d^+, c^+; d^-, c^-)$ in $\mathcal{D}(H)$, the procedure described in Section 8.1 selects a number d^+ of G^+ sites and a number d^- of G^- sites which determine a set of sites $S(d)$ of cardinal $d = d^+ + d^-$. The procedure given in Section 8.2 then selects a set $\mathcal{C}^+$ of c^+ cliques among pairs of G^+ biomarkers, and a set $\mathcal{C}^-$ of c^- cliques among pairs of G^- biomarkers.

On the set $\mathcal{B}(d)$ of binary vectors indexed by $S(d)$, we can then define as follows two autologistic models π^+ and π^- parameterized by the unknown parameter vectors θ^+ and θ^-

$$\pi^+ = \pi_{\theta^+}(x) = \frac{1}{Z(\theta^+)} e^{-<\theta^+, U(x)>} \quad and \quad \pi^- = \pi_{\theta^-}(x) = \frac{1}{Z(\theta^-)} e^{-<\theta^-, U(x)>},$$

with

$$< \theta^+, U(x) >= \sum_{s \in S(d)} \theta^+_s x_s + \sum_{\{s,t\} \in \mathcal{C}^+} \theta^+_{s,t} x_s x_t, \quad and \quad < \theta^-, U(x) >= \sum_{s \in S(d)} \theta^-_s x_s + \sum_{\{s,t\} \in \mathcal{C}^-} \theta^-_{s,t} x_s x_t.$$

To fit the models $\pi^+ = \pi_{\theta^+}$ and $\pi^- = \pi_{\theta^-}$ to the two data sets $BG^+ \subset \mathcal{B}(d)$ and $BG^- \subset \mathcal{B}(d)$, we estimate θ^+ and θ^- by the MPLE technique outlined in Section 5. Two separate gradient descent algorithms implement the computation of estimates $\hat{\theta}^+$ and $\hat{\theta}^-$. As explained in Sections 5.3 and 5.4, after computation of the 90% error margins on these estimated parameters, the coordinates of $\hat{\theta}^+$ and $\hat{\theta}^-$ which are not significantly different from 0 are then forced to be zero.

This autologistic model fitting procedure is repeated for each quadruplet $dim \in \mathcal{D}(H)$. This generates a number $N \sim O(H^2)$ pairs of autologistic models π^+ and π^-, all derived from the same original mass spectra data sets G^+ and G^-.

8.4 *Numerical Estimation of Optimal Separator*

Fix any quadruplet $dim = (d^+, c^+; d^-, c^-)$ in $\mathcal{D}(H)$. Let π^+ and π^- be the two autologistic models associated to dim and resp. fitted to the data sets BG^+ and BG^-.

Formula (7.10) for the optimal decision rule $f(x)$ discriminating between π^+ and π^- involves the term $\log(Z^-/Z^+)$. But computing the partition function $Z(\theta)$ of any autologistic π_θ by summing the $v(x) = e^{<\theta, U(x)>}$ over all $x \in \mathcal{B}(d)$ quickly becomes a heavy numerical task for $d > 16$. A known faster approach, justified by the law of large numbers, is to estimate $Z(\theta)/2^d$ by the average of $v(W_1), \ldots, v(W_N)$ where $N \sim 10,000$ and the W_j are independent random binary vectors generated by the *uniform* distribution on $\mathcal{B}(d)$. To avoid the errors involved in estimating $\log(Z^-/Z^+)$, and to partially compensate for parameters estimation errors, we have preferred to implement the following algorithm.

For each binary vector x in $\mathcal{B}(d)$, define the planar point $W(x)$ in R^2 by

$$W(x) = (w^+(x), w^-(x)) \quad with \quad w^+(x) =< \hat{\theta}^+, U(x) > \quad and \quad w^-(x) =< \hat{\theta}^-, U(x) > . \quad (8.16)$$

The non linear function $x \to W(x)$ thus transforms the subsets BG^+ and BG^- of $\mathcal{B}(d)$ into two sets $PLG^+ \subset R^2$ and $PLG^- \subset R^2$ of planar points labeled by $+1$ when $x \in BG^+$ and -1 when $x \in BG^+$. The optimal separator $f(x)$ defined by (7.10) is equivalent to an affine function $\mathcal{A}$ of

$W(x) \in R^2$ given by

$$\mathcal{A}.W(x) = f(x) = w^-(x) - w^+(x) + Constant. \tag{8.17}$$

Therefore, to generate a robust estimate $\hat{f}$ of the optimal separator f, we can directly search for an affine function $\mathcal{A}$ defined on the plane R^2, and which separates the two finite sets PLG^+ and PLG^- with a small number of errors.

To compensate for the errors due to the estimations of θ^+ and θ^-, which generate errors on the planar point $W(x)$, we introduce an error correcting coefficient β in the affine separator $\mathcal{A}$ as follows,

$$\mathcal{A}.W(x) = \hat{f}(x) = \beta w^-(x) - w^+(x) + Constant, \tag{8.18}$$

where β can be slightly different from 1. Given the two planar sets PLG^+ and PLG^- the "best" affine separator between PLG^+ and PLG^- can be quickly computed by many classical affine discrimination algorithm, such as Support Vector Machines with linear kernels (Fan *and others* (2008)). In each discrimination task G^+ vs G^-, the computation of the best separator $\hat{f}$ between π^+ and π^- has to be repeated for each quadruplet dim in $\mathcal{D}(H)$. So to gain in CPU time, we estimate $\mathcal{A}$ by linear regression of an indicator matrix (Hastie *and others* (2009)).

8.5 *Performance Evaluation for Autologistic Separator*

For each fixed quadruplet dim in $\mathcal{D}(H)$, we need to evaluate the probabilities $p^+(dim)$ and $p^-(dim)$ of successful discrimination between G^+ and G^- based on the "ideal" but unknown decision rule f. These probabilities can be evaluated by classical leave-one-out cross validation (Geisser (1993)).

Namely, at each cross validation round, one single binarized spectrum $x(M)$ is temporarily eliminated from $BG^+ \cup BG^-$, and this modified dataset is used to generate (as just outlined above) an estimated separator $F = \hat{f}$ of f. One checks then the sign of $F(x(M))$, to record whether the "left out" $x(M)$ is correctly classified by F or not. This procedure is repeated until every $x(M)$ in the dataset has been left out once. The *leave-one-out estimates* $\hat{p}^+$, $\hat{p}^-$ of $p^+(dim)$, $p^-(dim)$ are then the respective percentages of correct classifications of $x(M) \in BG^+$ and of $x(M) \in BG^-$. The *empirical performance* $PERF(dim)$ is then evaluated by $(\hat{p}^+ + \hat{p}^-)/2$.

The best choice dim_{opt} for the quadruplet $dim = (d^+, c^+; d^-, c^-)$ is then ideally determined

by maximizing the empirical performance $PERF(dim)$ over all dim in $\mathcal{D}(H)$. However, since the cardinal of $\mathcal{D}(H)$ ranges between 2,000 and 5,000 in all our benchmark studies, the full leave-one-out computation is too costly to be used for the computation of *all* the $PERF(dim)$ values. Fortunately, it is unnecessary to do so in numerical implementations, as will be seen in Section 8.7 below.

8.6 *Optimal Signature and associated Scores*

For the discrimination task G^+ vs G^-, once the best quadruplet $dim_{opt} = (d^+, c^+; d^-, c^-)$ has been determined by performance maximization, one has immediate access to the corresponding best autologistic models π^+_{opt} and π^-_{opt}, based on the selection of $d = d^+ + d^-$ specific reference biomarkers $B_{s(1)}, \ldots, B_{s(d)}$, involving numbers d^+ and d^- of G^+ and G^- biomarkers, as well as c^+ pairs of G^+ biomarkers and c^- pairs of G^- biomarkers.

The list $Sig = \{B_{s(1)}, \ldots, B_{s(d)}\}$ then constitutes an explicit optimized *signature* of length d to discriminate between G^+ and G^-.

Indeed let $\hat{f}_{opt}$ be the best separator between π^+_{opt} and π^-_{opt}, as computed in (7.11). For any mass spectrum M, let $X(M)$ be the binarized vector associated to M, and let $x(M)$ be the restriction of this binarized vector to the d signature sites $[s(1), \ldots, s(d)]$. Classifying the mass spectrum M into G^+ or into G^- requires only to know the sign of $F(M) = \hat{f}_{opt}(x(M))$. But the computation of $F(M))$ involves only checking the presence or absence in $X(M)$ of each one of the biomarkers belonging to the signature Sig.

Moreover, classification of mass spectra can then be interpreted as follows. Let $\hat{\theta}^+$ and $\hat{\theta}^-$ be the vectors parametrizing π^+_{opt} and π^-_{opt}. For any $j, k = 1 \ldots d$, define the *scores* $SCO(j)$ and $SCO(j,k)$

$$SCO(j) = \beta\hat{\theta}^-_{s(j)} - \hat{\theta}^+_{s(j)}, \quad SCO(j,k) = \beta\hat{\theta}^-_{s(j),s(k)} - \hat{\theta}^+_{s(j),s(k)},$$

where β was estimated in (8.18). For any mass spectrum M, we add up the scores $SCO(j)$ of all biomarker $B_{s(j)}$ which are present in M, as well as the scores $SCO(j,k)$ of all pairs of biomarker $B_{s(j)}, B_{s(k)}$ which are jointly present in M. Adding to this sum the constant occuring in formula (8.18) yields a *total score* $TSCO(M)$, which is actually equal to $F(M) = \hat{f}_{opt}(x(M))$, due to formulas (8.16) and (8.18).

Hence M is classified as belonging to G^+ or to G^- according to the sign of its total score

$TSCO(M)$.

8.7 *Accelerated Optimization of Discrimination Performances*

To estimate the best quadruplet dim_{opt} by performance maximization in realistic computing time, we have developed and implemented the following accelerated optimization scheme. Note first that in all our benchmark studies, the 90% error margins on $PERF(dim)$ are typically ~ 0.08.

For each quadruplet dim, once the autologistic models π^+, π^-, and their separator $\hat{f}$ have been estimated, one can instantaneously evaluate the empirical frequencies of good decisions achieved by $\hat{f}$ on the *training sets* G^+ and G^-. The average $\widehat{PERF}(dim)$ of these two frequencies provides a rough "training" approximation of $PERF(m)$. Once $\widehat{PERF}(dim)$ has been obtained for all dim, its maximization is immediate and provides first rough estimates $\hat{dim}_{opt}$ and $\widehat{PERF}_{opt}$ for the optimal quadruplet and the true best performance.

We can then compute $PERF(dim)$ by the costly leave-one-out technique, but only on the quite smaller subset $\mathcal{SD}(H)$ of all quadruplets dim which verify the constraint $\widehat{PERF}(dim) > \widehat{PERF}_{opt} - 0.08$. Maximizing $PERF(dim)$ for dim in $\mathcal{SD}(H)$ gives us our final estimate for dim_{opt} and the associated optimal performance.

For each one of our six benchmark studies, the computing time for best signature discovery was thus reduced to a range of 3 to 9 hours on a 1.3 GHz MacOS PC.

9. Kullback Distance Between Autologistic Models

Since the computation of performance $PERF(dim)$ is rather costly as seen above, we have also studied whether it was efficient to seek the best quadruplet dim_{opt} by maximizing over all $dim \in \mathcal{D}(H)$ the properly normalized Kullback-Leibler distance $norKL(dim)$ between the two autologistic models π^+ and π^- associated to each quadruplet $dim = (d^+, c^+; d^-, c^-)$.

Recall that the Kullback-Leibler (KL) divergence, introduced in Kullback and Leibler (1951), defines a well known distance non-negative distance $KL(P, Q)$ between probability distributions P and Q on a finite set $\mathcal{B}$ (Dagan *and others* (1999), Bigi *and others* (2003), Bigi (2003)), given

by

$$KL(P,Q) = \sum_{x \in \mathcal{B}} (P(x) - Q(x)) \log(\frac{P(x)}{Q(x)}).$$

and verifying $KL(P,Q) = 0$ iff $P = Q$.

For autologistic models $\pi^+ = \pi_{\theta^+}$ and $\pi^- = \pi_{\theta^-}$, we thus have

$$KL(\pi^+, \pi^-) =< \theta^- - \theta^+, \sum_{x \in \mathcal{B}} (\pi^+(x) - \pi^-(x))U(x) > . \tag{9.19}$$

Recall that $f(x) =< \theta^- - \theta^+, U(x) > + \log(Z^-/Z+)$ is the optimal separator between π^+ and π^-, we hence have

$$KL(\pi^+, \pi^-) = \sum_{x \in \mathcal{B}} f(x)\pi^+(x) - \sum_{x \in \mathcal{B}} f(x)\pi^-(x).$$

When the optimal separators f has high performance, $f(x)$ must be positive with high probability under π^+ and negative with high probability under π^-. In view of the last formula, we should then expect the distance $KL(\pi^+, \pi^-)$ to typically be high as well.

This last statement requires a proper normalization of KL distances, to take into account the dimension $d = d^+ + d^-$ of the set of binary vectors $\mathcal{B}(d)$, since the autologistic probability distributions π^+ and π^- live on $\mathcal{B}(d)$. Indeed, a brief analysis of the case where pairs of sites are very weakly correlated suggests to normalize KL distances as follows

$$norKL(\pi^+, \pi^-) = KL(\pi^+, \pi^-)/\sqrt{d^+ + d^-}, \tag{9.20}$$

and for each dimension quadruplet dim, we will then denote $norKL(dim) = norKL(\pi^+, \pi^-)$.

Our MRF discrimination algorithm searches for "well separated" pairs π^+ and π^- by maximizing performance $PERF(dim)$ over all dimensions quadruplets dim. We conjecture that one can achieve roughly the same goal by maximizing $norKL(dim)$ over all dim. Since $norKL(dim)$ is much easier to compute than $PERF(dim)$, the maximization of $norKL(dim)$ seems to be a rougher but faster approach to the discovery of the best quadruplet dim_{opt}.

As we have numerically checked in our six benchmark studies, this conjecture is correct at the accuracy level 0.08 of all our performance evaluations. Indeed, our numerical results below show that normalized KL distances $norKL(dim)$ are positively correlated to the actual discrimination performances $PERF(dim)$, and that this correlation is stronger at high values of $PERF(dim)$.

10. Application to Benchmark Cancer Datasets

10.1 *Pre-Processing of Benchmark Mass Spectra*

For the two ovarian cancer datasets, pre-processing had practically already been performed in Kong *and others* (2014) and provided us with a fixed list of 15,154 m/z ratios ranging from 0 to 20,000 and which thus defined our list of 15,154 reference biomarkers.

For the colorectal cancer dataset, we did implement the pre-processing of raw mass spectra as detailed in our joint paper Kong *and others* (2014), and this generated a list of 842 reference biomarkers positioned at the 842 m/z ratios $800 \times (1+\rho)^j$ with $\rho = 0.3\%$ and $j = 0, 1, 2, \ldots, 841$.

As indicated in Section 3.2, each pre-processed mass spectrum M was then coded as a binary vector $x(M)$ of length 15,154 for ovarian cancer data, and of length 842 for colorectal cancer data.

10.2 *Benchmark Implementations of MRF Discrimination*

For each one of the 6 discrimination tasks outlined in Section 2, generically denoted G^+ vs G^-, we have systematically explored all the dimensions quadruplets $dim = (d^+, c^+; d^-, c^-)$ belonging to $\mathcal{D}(H)$ where $H = 15$ for the three colorectal cancer stage discrimination tasks ADE vs ECR, ADE vs LCR, ECR vs LCR, and $H = 10$ for the three cancer vs control discrimination tasks on both colorectal and ovarian cancer data. These deliberate restrictions forced two moderate a priori upper bounds for the numbers $d = d^+ + d^-$ of biomarkers explored in signature discovery, namely signature lengths were thus kept inferior to 30 for colorectal cancer data, and to 20 for ovarian cancer data. Besides the associated reductions in computing time, these bounds also forced the number of signature scores to remain inferior to $n/5$ where n was the total number of mass spectra in the training set $G^+ \cup G^-$.

For each quadruplet $dim \in \mathcal{D}(H)$, we have implemented our MRF discrimination algorithm to the pair G^+, G^- of binarized mass spectra, as outlined in Section 8.4. This yielded two estimated autologistic models π^+, π^- for G^+, G^- and the associated best separator $\hat{f}$ between π^+ and π^-. The frequencies p^+ and p^- of correct decisions achieved by $\hat{f}$ on the training sets G^+ and G^- were then directly evaluated to derive a first rough estimate $\widehat{PERF}(dim) = (p^+ + p^-)/2$ for the "training performance" of the separator $\hat{f}$. After completing this computation of $\widehat{PERF}(dim)$ for

all $dim \in \mathcal{D}(H)$, the accelerated procedure outlined in Section 8.7 then enabled the computation of the optimal quadruplet $dim_{opt} = (d^+, c^+; d^-, c^-)$ maximizing $PERF(dim)$ and the associated separator $\hat{f}_{opt}$ discriminating between G^+ and G^-.

As indicated in Section 8.6, this determined an optimized signature $Sig = \{B_{s(1)}, \ldots, B_{s(d)}\}$ gathering $d = d^+ + d^-$ specific reference biomarkers to discriminate G^+ vs G^-. This computation also yielded numerical scores $SCO(i)$ and $SCO(i,j)$ associated to each signature biomarker $B_{s(i)}$ and to each retained clique $\{B_{s(i)}, B_{s(j)}\}$. As seen in Section 8.6, classifying a mass spectrum M into G^+ or G^- depends then only on the sign of its total score $TSCO(M)$, which can be evaluated by checking the presence/absence in M of the signature biomarkers.

10.3 *Benchmark Performances of MRF Discrimination*

As seen above, once the two optimized autologistic models π^+, π^- have been estimated, the mass spectra datasets G^+, G^- can be mapped into two sets PLG^+ and PLG^- of planar points by the non linear map $M \to x = x(M) \to W(x)$, and the optimized classifier $\hat{f}_{opt}(x)$ is of the form $\mathcal{A}(W(x)$ where $\mathcal{A}$ is an affine function on R^2 which separates PLG^+ and PLG^- with maximum margin. For each benchmark discrimination task on the colorectal cancer dataset, we have displayed PLG^+, PLG^- and the associated best affine separator $\mathcal{A}$ in Figure 2.

Our MRF discrimination algorithm thus discovered 6 explicit optimized signatures, one for each benchmark discrimination task. We first note that these 6 signatures all had fairly short lengths $d = d^+ + d^-$, namely $d = 23, 17, 18, 7$ for the 4 colorectal cancer discrimination tasks, and $d = 12, 10$ for the two ovarian cancer discrimination tasks. These short signature lengths obviously are a strong advantage for further biological interpretation of these 6 small families of key biomarkers, explicitly identified by their m/z ratios.

The performances achieved on our 6 benchmark discrimination tasks by the optimal separators associated to these 6 signatures are displayed in Table 1. Taking into account the error margins on estimates of p^+, p^-, our MRF discrimination algorithm reached performance levels $perf(MRF) = (p^+ + p^-)/2$ which were *essentially equivalent* to the best performances reported by previous publications, which used other discrimination algorithms.

Indeed, on colorectal cancer data the discrimination performances $perf(SA)$ reached by intensive Simulated Annealing for signature discovery (Kong *and others* (2014)) and the performances

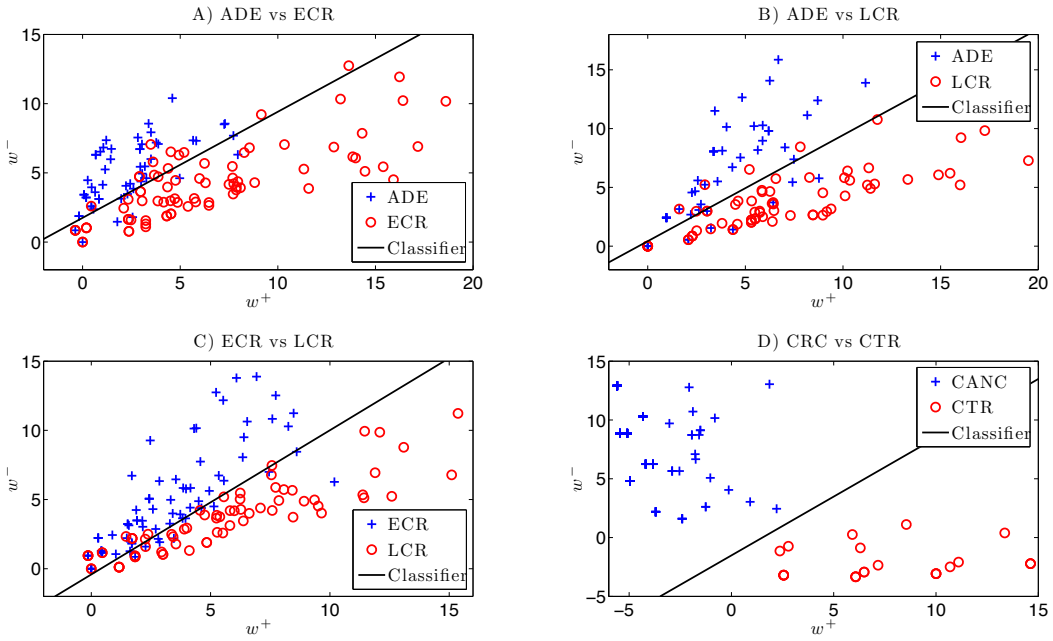

Fig. 2. Colorectal Cancer Data: For each one of four discrimination tasks G^+ vs G^- we display the planar representation of these two groups of mass spectra, derived from the optimized autologistic models fitted to the data sets G^+ and G^-. The associated non linear recoding computes two planar coordinates (w^+, w^-) for each mass spectrum. The best non linear separator between G^+ and G^- can then be displayed as an affine function of (w^+, w^-).

Table 1. *Performance of optimized MRF Discrimination between G^+ and G^- is evaluated by the frequencies p^+ and p^- of correct classifications within G^+ and G^-. For each discrimination task, we display p^+, p^- with their 90% confidence intervals and the quadruplet $(d^+, c^+; d^-, c^-)$ of dimension parameters for the optimal autologistic separator between G^+ and G^-, which is based on a "signature" involving only $d = d^+ + d^-$ specific reference biomarkers.*

	COLORECTAL CANCER			
	ADE vs ECR	ADE vs LCR	ECR vs LCR	CRC vs CTR
$(d^+, c^+; d^-, c^-)$	(14, 4; 9, 1)	(4, 1; 13, 1)	(7, 0; 11, 2)	(2, 0; 5, 0)
p^+ vs p^-	0.74 vs 0.88	0.81 vs 0.92	0.83 vs 0.82	0.995 vs 1
90% conf. int.	[0.64; 0.84] [0.82; 0.94]	[0.72; 0.90] [0.87; 0.97]	[0.76; 0.90] [0.75; 0.89]	[0.99; 1] [1; 1]

	OVARIAN CANCER (4-3-02) Cancer vs Control	OVARIAN CANCER (8-7-02) Cancer vs Control
$(d^+, c^+; d^-, c^-)$	(7, 2; 5, 1)	(6, 0; 4, 0)
p^+ vs p^-	0.93 vs 0.91	1 vs 0.98
90% conf. int.	[0.89; 0.97] [0.87; 0.95]	[1; 1] [0.96; 1]

$perf(MRF)$ reached in the present paper were the following, with error margins of the order of 0.08:

for ADE vs ECR: $perf(MRF) = 0.81, perf(SA) = 0.84$,

for ADE vs LCR: $perf(MRF) = perf(SA) = 0.86$,

for ECR vs LCR: $perf(MRF) = 0.82, perf(SA) = 0.83$,

for CRC vs CTR: $perf(MRF) = 0.99, perf(SA) = 1$.

For discrimination task Cancer vs Control on the ovarian cancer dataset 4-3-02, our MRF approach reached the performance level $perf(MRF) = 0.92$, while the papers Assareh and Moradi (2007) and Kong *and others* (2014) resp. reported performances 0.92 and 0.94.

For the easier discrimination of Cancer vs Control on the Ovarian Cancer dataset 8-7-02, our MRF approach reached $perf(MRF) = 0.99$, which is equivalent to the performances of 100% reported by the four papers Zhu *and others* (2003), Assareh and Moradi (2007), Alexe *and others* (2004), Kong *and others* (2014).

10.4 *MRF Discrimination Performance and KL Distance*

For each benchmark discrimination task G^+ vs G^-, and for each quadruplet $dim = (d^+, c^+; d^-, c^-)$ in $\mathcal{D}(H)$, where $H = 15$ or 10, we have fitted as above two autologistic models π^+ and π^- to the data sets G^+ and G^- and computed the best separator between these two models. After evaluating the discrimination performance $PERF(dim)$ of this separator, we have also systematically computed the normalized KL distance $norKL(dim)$ between π^+ and π^- (see Section 9). To roughly evaluate the correlation between $PERF(dim)$ and $norKL(dim)$ we have then plotted the planar points $(PERF(dim), norKL(dim))$ on R^2. An example of such a plot is given in Figure 3 for the discrimination task ECR vs LCR, based on colorectal cancer data. We have observed a fuzzy but positive correlation between autologistic discrimination performance $PERF$ and normalized KL distance $norKL$ between pairs autologistic models fitted to G^+ and G^-. We have also noted numerically that maximization of $PERF(dim)$ is roughly equivalent to maximization of $norKL(dim)$ over all dimension quadruplets dim in $\mathcal{D}(H)$.

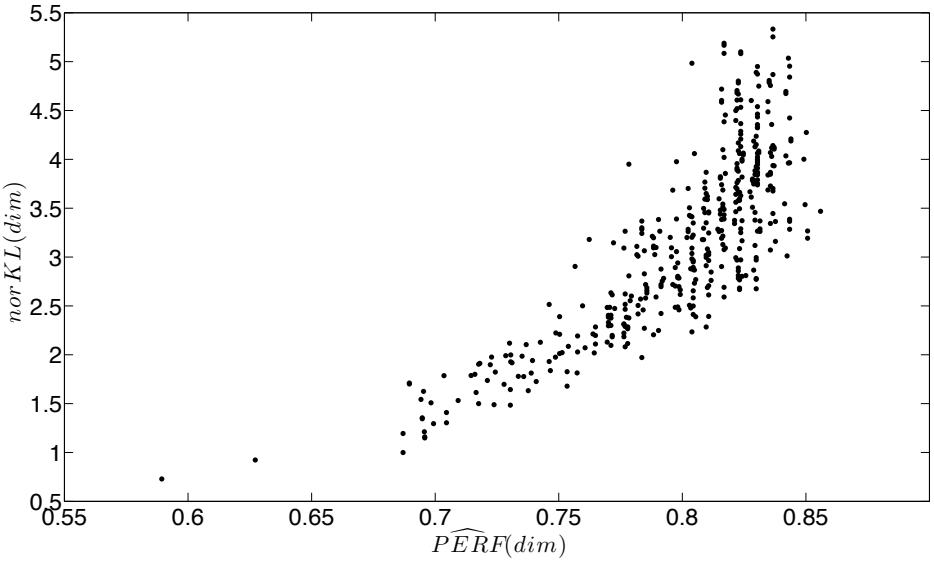

Fig. 3. After fitting 2538 pairs π_j^+, π_j^- of autologistic models to the two groups $G^+ = ECR$ and $G^- = LCR$ we have computed the performance $PERF_j$ of the best separator between π_j^+ and π_j^-, as well as the normalized KL distance $norKL_j$ between π_j^+ and π_j^-. The figure displays the 2538 planar points with coordinates $(PERF_j, norKL_j)$. The graph indicates a positive correlation between $PERF_j$ and $norKL_j$ particularly for high performances.

.

11. Detailed Implementation for Benchmark Discrimination Task ECR vs LCR

11.1 *Sparse Selection of ECR-biomarkers and LCR-biomarkers*

As just reported, our MRF discrimination approach was successfully tested on cancer data for 6 benchmark discrimination tasks. We now sketch more detailed implementation steps, but only for one example, namely the discrimination of $G^+ = ECR$ versus $G^- = LCR$. This data set involved 154 mass spectra: 80 for G^+ patients and 74 for G^- patients. Pre-processing of these 154 spectra (see Section 3.1) yielded a list of 842 reference biomarkers B_s, and hence an initial set S of 842 sites s. Binary coding (see Section 3.2) of mass spectra then generated two sets $BinG^+$ and $BinG^-$ of binary vectors of length 842, with respective cardinals 80 and 74.

For each $s \in S$, let $m^+(s)$ and $m^-(s)$ be the frequencies of the event $x_s = 1$ within $BinG^+$ and $BinG^-$ respectively. We then focus only on the set of sites $\hat{S}$ of all $s \in S$ such that $m^+(s) \geqslant 20\%$ and $m^-(s) \geqslant 20\%$, and as above we let $DP(s) = m^+(s)/m^-(s)$.

For each such quadruplet $dim \in \mathcal{D}(15)$, we select as G^+ and G^- biomarkers the d^+ sites $s \in \hat{S}$ with highest $DP(s)$ and the d^- sites s with lowest $DP(s)$. As in Section 8.2, we then select c^+ pairs of G^+ biomarkers, and c^- pairs of G^- biomarkers. The binarized spectra $BinG^+$ and $BinG^-$ are then restricted to the $d = d^+ + d^-$ sites just selected to define two sets BG^+ and BG^- of binary vectors of length d.

11.2 *ECR vs LCR: Model fitting and Maximization of Performance*

As in Section 8.3, we fit two autologistic distributions π^+ and π^- to BG^+ and BG^-, to then compute an optimized separator $\hat{f}$ between π^+ and π^-. The empirical performance $\widehat{PERF}(dim)$ of $\hat{f}$ is roughly estimated as the average percentage of correct discrimination decisions generated by $\hat{f}$ on the training set.

The computation of $\widehat{PERF}(dim)$ is repeated for all $dim \in \mathcal{D}(15)$, and maximization of $PERF(dim)$ over all dim by the accelerated procedure of Section 8.7 yields the best quadruplet dim_{opt}. For the case $G^+ = ECR$ vs $G^- = LCR$, this dim_{opt} was $d^+ = 7, c^+ = 0; d^- = 11, c^- = 2$, achieving maximal performance $PERF_{opt} = 0.82$. The associated separator $\hat{f}_{opt}$ reached percentages of good decisions $p^+ = 0.83$ on ECR data and $p^- = 0.82$ on LCR data.

The best autologistic models π^+_{opt} and π^-_{opt} associated to dim_{opt} are based on a signature of length 18, involving 7 G^+ biomarkers $B_1, \ldots, B_7$ and 11 G^- biomarkers $B_8, \ldots, B_{18}$, identified by their m/z ratios (in Daltons) listed in Table 2. No clique of G^+ biomarkers was retained for π^+, but π^- involved 2 cliques of G^- biomarkers which are listed as well as their χ^2-statistics in Table 3. The parameter vectors $\hat{\theta}^+$ and $\hat{\theta}^-$ of π^+_{opt} and π^-_{opt} generated by MPLE are of dimension 20, and are displayed in two columns in Table 4, which also gives the error margins on parameter estimates. These error margins define the 90% confidence intervals computed from asymptotic normality results (Section 5.3). The zero values displayed for $\hat{\theta}^+_3$, $\hat{\theta}^+_4$, $\hat{\theta}^+_6$, $\hat{\theta}^-_{12}$, $\hat{\theta}^-_{18}$ were imposed after a first MPLE evaluation and error margins computation indicating that these parameters were not significantly different from 0 at the 90% confidence level.

The quality of fit of π^+ to ECR data and of π^- to LCR data were evaluated as in Section 5.5, and LCR data and yielded good respective quantile values 46.8% and 51.3% for the log-likelihoods of the observed ECR and LCR data.

Table 2. *This table presents the 18 biomarkers retained for optimized discrimination between colorectal cancer groups $G^+ = ECR$ and $G^- = LCR$, as well as their m/z ratios (in Daltons). The columns $m^+(s)$ and $m^-(s)$ display the respective frequencies of biomarker B_s activation by the mass spectra in G^+ and G^-. The top and bottom panels respectively present the 7 biomarkers with the highest and $m^+(s)/m^-(s)$ ratios and the 11 biomarkers with the highest $m^-(s)/m^+(s)$ ratios.*

G^+ biomarkers				
B_s	m/z	$m^+(s)$	$m^-(s)$	$m^-(s)/m^+(s)$
B_1	1953	0.21	0.06	3.22
B_2	1112	0.29	0.14	2.09
B_3	8376	0.55	0.31	1.78
B_4	6320	0.43	0.25	1.76
B_5	2013	0.24	0.14	1.76
B_6	921	0.44	0.26	1.72
B_7	2558	0.30	0.17	1.71
G^- biomarkers				
B_s	m/z	$m^+(s)$	$m^-(s)$	$m^-(s)/m^+(s)$
B_8	7889	0.09	0.26	2.77
B_9	2807	0.15	0.34	2.29
B_{10}	1032	0.13	0.26	2.03
B_{11}	819	0.17	0.32	1.85
B_{12}	2143	0.26	0.48	1.81
B_{13}	1930	0.12	0.21	1.81
B_{14}	2269	0.12	0.21	1.81
B_{15}	2815	0.12	0.21	1.81
B_{16}	2331	0.17	0.31	1.78
B_{17}	1478	0.15	0.26	1.72
B_{18}	1210	0.26	0.45	1.72

Table 3. *This table lists the cliques of order 2 for the best autologistic models π^+ and π^- in the discrimination task $G^+ = ECR$ vs $G^- = LCR$. No cliques of G^+ biomarkers were retained for π^+, and 2 cliques of G^- biomarkers were selected for π^-. The principle is to retain cliques with highest χ^2-statistics. The table displays the 2 selected cliques and the corresponding χ^2-statistics.*

clique $\{s,t\}$	m/z of B_s	m/z of B_t	$\chi^2(s,t)$
{9, 10}	2807	1032	7.15
{13,18}	1930	1210	5.45

11.3 *ECR vs LCR: Optimized Discriminating Signature and Scores*

The optimized signature discovered for ECR vs LCR discrimination involves 7 ECR-biomarkers $B_1, \ldots, B_7$ and 11 LCR-biomarkers $B_8, \ldots, B_{18}$, as well as the 2 pairs of LCR-biomarkers (B_9, B_{10}) and (B_{13}, B_{18}). As indicated in Section 8.6, automatic classification of a mass spectrum M into either ECR or LCR is based on the 18 scores $SCO_1, \ldots, SCO_{18}$ and the 2 scores $SCO(9,10)$, $SCO(13,18)$. These scores are the coordinates of $\beta\hat{\theta}^- - \hat{\theta}^+$ and the computed value for β was

0.92. These 20 scores are listed in Table 4.

Table 4. *This table lists the G^+ and G^- biomarkers as well as the only 2 cliques of order 2 retained for the two autologistic models π^+ and π^- enabling optimal discrimination between $G^+ = ECR$ and $G^- = LCR$. The columns $\hat{\theta}^+$ and $\hat{\theta}^-$ display the parameter vectors of π^+ and π^- as estimated by MPLE. For each coordinate of $\hat{\theta}^+$ and $\hat{\theta}^-$, we also give the 90% confidence intervals derived from the asymptotic normality results. The zero values displayed for $\hat{\theta}^+_3$, $\hat{\theta}^+_4$, $\hat{\theta}^+_6$, $\hat{\theta}^-_{12}$, $\hat{\theta}^-_{18}$ were imposed after a first estimation indicating that these parameters were not significantly different from 0. The two cliques $\{s,t\}$ involve only G^- biomarkers so that the corresponding $\hat{\theta}^+_{s,t}$ are equal to zero. The table also lists the 20 discrimination scores $SCO(s)$ and $SCO(s,t)$ associated to this signature.*

G^+ biomarkers					
B_s	$\hat{\theta}^+_s$	conf. int. of $\hat{\theta}^+_s$	$\hat{\theta}^-_s$	conf. int. of $\hat{\theta}^-_s$	$SCO(s)$
B_1	1.6	[1.1; 2.1]	4.3	[2.5; 6.0]	2.4
B_2	1.0	[0.6; 1.4]	2.3	[1.6; 3.0]	1.0
B_3	0	-	0.9	[0.5; 1.3]	0.9
B_4	0	-	1.3	[0.8; 1.8]	1.2
B_5	1.3	[0.8; 1.8]	2.3	[1.6; 3.0]	0.8
B_6	0	-	1.2	[0.7; 1.7]	1.1
B_7	1.0	[0.6; 1.4]	1.9	[1.3; 2.5]	0.7
G^- biomarkers					
B_s	$\hat{\theta}^+_s$	conf. int. of $\hat{\theta}^+_s$	$\hat{\theta}^-_s$	conf. int. of $\hat{\theta}^-_s$	$SCO(s)$
B_8	2.9	[2.1; 3.7]	1.2	[0.7; 1.7]	-1.8
B_9	2.1	[1.5; 2.7]	1.1	[0.6; 1.6]	-1.0
B_{10}	2.3	[1.6; 3.0]	1.8	[1.1; 2.5]	-0.7
B_{11}	1.8	[1.3; 2.3]	0.9	[0.5; 1.3]	-1.0
B_{12}	1.2	[0.8; 1.6]	0	-	-1.2
B_{13}	2.5	[1.8; 3.2]	1.0	[0.4; 1.6]	-1.6
B_{14}	2.5	[1.8; 3.2]	1.5	[1.0; 2.0]	-1.1
B_{15}	2.5	[1.8; 3.2]	1.5	[1.0; 2.0]	-1.1
B_{16}	1.8	[1.3; 2.3]	0.9	[0.5; 1.3]	-1.0
B_{17}	2.1	[1.5; 2.7]	1.2	[0.7; 1.7]	-1.0
B_{18}	1.2	[0.8; 1.6]	0	-	-1.2
G^- cliques					
$\{s,t\}$	$\hat{\theta}^+_{s,t}$	conf. int. of $\hat{\theta}^+_{s,t}$	$\hat{\theta}^-_{s,t}$	conf. int. of $\hat{\theta}^-_{s,t}$	$SCO(s,t)$
{9, 10}	0	-	-1.5	[-2.5; -0.5]	-1.4
{13, 18}	0	-	1.7	[0.4; 3.0]	1.6

For any mass spectrum M and any $s = 1, \ldots, 18$, recall that we have set $x_s(M) = 1$ when the biomarker B_s is activated by M, and $x_s(M) = 0$ otherwise. The total score $TSCO(M)$ is here given by

$$TSCO(M) = 0.37 + x_1(M)SCO(1) + \ldots + x_{18}(M)SCO(18) + x_9(M)x_{10}(M)SCO(9,10) + x_{13}(M)x_{18}(M)SCO(13,18)$$

Then M is classified as belonging to the cancer group ECR if $TSCO(M) > 0$ and to the group

LCR otherwise.

12. Conclusion

Mass spectrometry is a promising approach for biomarker-based early cancer detection. By algorithmic analysis of mass spectra acquired by MALDI-TOF or SELDI-TOF techniques, selected sets of peptides strongly linked to specific cancer groups have been used for automated cancer stage classification. To further help incorporating these approaches in reliable clinical protocols, a key step is to discover interpretable "biomarker signatures "characterizing various cancer types and/or cancer stages.

Machine learning algorithms, such as decision trees, support vector machines, artificial neural networks, etc., have been tested on mass spectra data sets acquired from cancer patients, and have often been proposed as efficient methods to discriminate between groups of mass spectra. But these techniques tend to generate "black-box" results which often lack direct biological interpretability.

Our main focus was to rigorously fit parameterized stochastic models to mass spectra datasets acquired by MALDI or SELDI techniques, in order to design efficient signature discovery algorithms leading to interpretable signatures combining the discriminating power of well selected *small* groups of biomarkers.

In this paper, the pattern variations observable in any supposedly "homogeneous " data set acquired by mass spectrometry have been systematically modeled by binary Markov Random Fields (MRF). After fairly classical pre-processing of any given group G of mass spectra, we generate a reference list of several hundreds to several thousands of strong spectral "peaks", viewed as potential key biomarkers to characterize the group G. We then code each mass spectrum M by a long binary vector listing the binary status (presence/absence in M) of each reference biomarker. Gibbs distributions, as used in our work, are efficient stochastic models to study the spatial dependency of coordinates for high dimensional binary vectors viewed as realizations of Markov Random Fields. We have focused our study on autologistic models - a type of Gibbs distributions involving only paiwise interactions between binary sites. Automatic classification of new mass spectra into two distinct mass spectra datasets G^+ and G^- can then be reduced to computing the best classifier discriminating two autologistic models.

Based on this theoretical point of view and associated algorithms, we have proposed a systematic approach to discover explicit and highly discriminating biomarker signatures, enabling efficient discrimination between distinct homogeneous groups of binary coded mass spectra. To construct autologistic models of random peaks patterns variations among mass spectra, stochastic model fitting to data was implemented by Maximum Pseudo-Likelihood Estimation (MPLE), and achieved a good quality of fit to mass spectra data in all our cancer data sets studies.

We have successfully tested our innovative signature discovery algorithms on a new experimental set of MALDI-TOF mass spectra acquired from patients at three stages of colorectal cancer, as well as on two previously published data sets of SELDI-TOF mass spectra acquired from ovarian cancer patients. Final performance levels are computed by leave-one-out cross validation. The performance of our algorithm is good in all these concrete cases and compared quite favorably with performance levels reported in previous publications. The clear concrete advantage of our optimized signature discovery technique is that it provides a biologically interpretable signature, involving only a small number of key biomarkers identified by their m/z ratios, and where each biomarker is weighted by an explicit numerical score.

Acknowledgments

The authors thank M. Agostini and C. Bedin at the 1st Surgical Clinic, Dept of Surgical, Oncological and Gastroenterological Sciences (University of Padova) for providing plasma samples from colorectal cancer patients; we also thank A. Bouamrani, E. Tasciotti and M. Ferrari from The Methodist Hospital Research Institute of Houston (Dept. of Nanomedicine) for providing us with the MALDI-TOF mass spectra acquired from these plasma samples. We have also benefitted from multiple interesting and informative discussions with the five specialists just mentioned, as well as with C. Gupta (Mathematics, University of Houston), in the framework of the collaboration which led to the joint publication Kong *and others* (2014).

Conflict of Interest: None declared.

Appendix

This mathematical appendix outlines concise and self-contained proofs of several theoretical results stated and used in this paper. Indeed, in the literature on MRFs, Gibbs distributions, and MPLE estimation on finite configuration spaces, similar results have been used at least implicitly, but elements of proof are often absent, or linked to very different contexts, or imbedded within proofs of other theorems.

In all the following theorems and proofs, $\mathcal{B}$ is a space of binary vectors of fixed length L.

A. Existence and Uniqueness of Maximum Pseudo-Likelihood Estimators

Consider as in Section 4, and with the same notations, an autologistic distribution π_θ parameterized by the unknown vector $\theta \in \Theta$. Given n observed random binary vectors generated by π_θ, the MPLE estimator $\hat{\theta}_n$ of θ is computed by maximizing in $T \in \Theta$ the empirical pseudo-likelihood $\hat{g}(T, n)$ as defined by (5.7). Existence and fast computability of the MPLE is due to the concavity of $\hat{g}(T, n)$ and of its limit $g_\theta(T)$ as $n \to \infty$, as we now prove.

Proposition 1 For each $\theta \in \Theta$, the mean log pseudo-likelihood $g_\theta(T)$ defined by equation (5.6) is a strictly concave function of $T \in \Theta$, and reaches its maximum in T at the unique point $T = \theta$. Moreover, the empirical pseudo-likelihood $\hat{g}(T, n)$ is also concave in $T \in \Theta$, and becomes almost surely strictly concave as $n \to \infty$.

Proof: (Notations of Section 5.1 and 5.2). The functions $A(s,x), g(z), h(z), LPL_{s,x}(T)$ were defined by equation (5.3), (5.4), and (5.5). For any binary number x_s, the function $\phi(z) = x_s\, z + \log(g(z))$ then verifies

$$\phi'(z) = x_s - h(z)\,; \qquad \phi''(z) = -e^z/(1+e^z)^2 < 0.$$

Since $LPL_{s,x}(T) = \phi(< A(s,x), T >)$, the gradient in T of $LPL_{s,x}(T)$ is the vector

$$\partial_T LPL_{s,x}(T) = \phi'(< A(s,x), T >)\, A(s,x) = [\, x_s - h(< A(s,x), T >)\,]\, A(s,x). \tag{A.21}$$

The Hessian $HLPL(s,x,T)$ of $LPL_{s,x}(T)$ with resp. to T is hence the quadratic form defined on all vectors $t \in \Theta$ by

$$t^*\, HLPL(s,x,T)\, t = \phi''(< A(s,x), T >)\ < A(s,x), t >^2$$

so that

$$t^*\, HLPL(s,x,T)\, t = \frac{- < A(s,x), t >^2 \;\; e^{<A(s,x),T>}}{(1+e^{<A(s,x),T>})^2} \leqslant 0.$$

The Hessian $Hg_\theta(T)$ of $g_\theta(T)$ with respect to T is hence the quadratic form defined on all vectors $t \in \Theta$ by

$$t^*\, Hg_\theta(T)\, t = \sum_s \sum_x \left[\, \pi_\theta(x) \phi''(< A(s,x), T >) < A(s,x), t >^2 \,\right] \leqslant 0.$$

The quadratic form $t^*\, Hg_\theta(T)\, t$ in $t \in \Theta$ is thus clearly non positive. Since $\pi_\theta(x) > 0$ and $\phi''(z) < 0$, this quadratic form takes the value 0 for some $t \in \Theta$ if and only if $< A(s,x), t >= 0$ for all $s \in S$ and all binary configurations x. But the conditions $t_s + \sum_{s<r} t_{sr} x_r = 0$ for all indices s and all binary vectors x are easily shown to force $t = 0$.

The Hessian of $g_\theta(T)$ is hence negative definite for all T, and thus the function $g_\theta(T)$ is strictly concave for all $T \in R^k$. Since g_θ is obviously bounded, it has a unique maximum in T, and this maximum is reached at the unique vector $\hat{T}$ which cancels the gradient $\partial_T g_\theta(T)$ of $g_\theta(T)$. So to conclude that $\hat{T} = \theta$, we now prove that $\partial_T g_\theta(\theta) = 0$.

Due to equations (A.21) and (5.2), for all indices s and $s < r$, and for all binary configurations x, the partial derivatives of $LPL(x,T)$ are given by

$$\frac{\partial LPL(x,T)}{\partial T_s} = -x_s + h(< A(s,x), \theta >), \tag{A.22}$$

$$\frac{\partial LPL(x,T)}{\partial T_{sr}} = -2x_s x_r + x_r\, h(< A(s,x), \theta >) + x_s\, h(< A(s,x), \theta >). \tag{A.23}$$

Letting $T = \theta$ in the conditional specifications formulas given above, we get, for a random configuration X with distribution π_θ,

$$P_\theta(X_s = 1 \mid X_{S-s}) = h(< A(s,X), \theta >)$$

and hence for all indices s and $s < r$ one has

$$E_\theta(h(< A(s,X), \theta >)) = E_\theta(P_\theta(X_s = 1 \mid X_{S-s}) = P_\theta(X_s = 1) = E_\theta(X_s),$$

$$E_\theta(X_r\, h(< A(s,X), \theta >)) = E_\theta(X_r P_\theta(X_s = 1 \mid X_{S-s})) = E_\theta(X_r X_s).$$

Due to equations (A.22), (A.23), these identities prove that for all indices s and $s < r$

$$\frac{\partial g_\theta(\theta)}{\partial T_s} = E_\theta \left(\frac{\partial LPL(X,\theta)}{\partial T_s} \right) = 0 \quad \textit{and} \quad \frac{\partial g_\theta(\theta)}{\partial T_{sr}} = E_\theta \left(\frac{\partial LPL(X,\theta)}{\partial T_{sr}} \right) = 0.$$

We have thus shown that the gradient $\partial_T g_\theta(T)$ is equal to 0 when $T = \theta$. This concludes the proof.

B. Asymptotic Normality of MPLE

Proposition 2 For any autologistic distribution π_θ on the finite set of binary configurations $\mathcal{B}$, the MPLE estimators $\hat{\theta}(n)$ of θ are asymptotically consistent as the number of observations $n \to \infty$. The normalized vectors of estimation errors $\sqrt{n}\ (\hat{\theta}(n) - \theta)$ are asymptotically Gaussian with mean zero and covariance matrix

$$\Gamma(\theta) = \mathcal{H}(\theta)^{-1}\Sigma(\theta)\mathcal{H}(\theta)^{-1},$$

where $\mathcal{H}(\theta)$ is the Hessian of $g_\theta(T)$ at $T = \theta$, and the symmetric positive definite matrix $\Sigma(\theta)$ is determined by equations (B.24) (see below).

Proof: (Notations of Section 5.1 and 5.2). We only outline the main technical steps, since our approach is similar to the "contrast function" analysis applied in the last chapter of Azencott and Dacunha-Castelle (1985).

For all fixed parameter vectors $\theta \in \Theta$ the contrast function $g_\theta(T)$ is strictly concave in T, as seen above, and the law of large numbers gives the almost sure limit $\lim_{n\to\infty} \hat{g}(T, n) = g_\theta(T)$.

Asymptotic consistency of $\hat{\theta}(n)$ is then derived as in Azencott and Dacunha-Castelle (1985).

Let $\hat{G}(T, n)$ and $\hat{H}(T, n)$ be the gradient and the Hessian of $\hat{g}(T, n)$ as defined by equations (5.8) and (5.9). Apply first the central limit theorem to assert that the random vector

$$Y(T, n) = \sqrt{n}\ (\hat{G}(T, n) - \partial_T g_\theta(T))$$

is asymptotically Gaussian with mean zero and covariance matrix

$$\Sigma(T) = E_\theta(\partial_T LPL(X, T)\ \partial_T LPL(X, T)^*), \tag{B.24}$$

where X is a random configuration with distribution π_θ.

Let $\Sigma_{s,j}\,, \Sigma_{st,ij}\,, \Sigma_{s,ij}$, with $s, t, i, j \in S$ and $s < t,\ i < j$, denote the elements of the symmetric matrix $\Sigma(T)$.

Let $W_s = h(< A(s, X), T >)$, and apply formulas (A.22) and (A.23) to obtain

$$\Sigma_{s,i}(T) = E_\theta[\ (-X_s + W_s)(-X_i + W_i)\ ],$$

$$\Sigma_{st,ij}(T) = E_\theta[\ (-2X_sX_t + X_tW_s + X_sW_t)(-2X_iX_j + X_jW_i + X_iW_j)\ ],$$

$$\Sigma_{s,ij}(T) = E_\theta[\ (-X_s + W_s)(-2X_iX_j + X_jW_i + X_iW_j)\ ].$$

For any given T, these expressions can easily be used to compute good numerical estimates of $\Sigma(T)$ by simulated Gibbs sampling of π_θ.

For T close to θ, and since $\partial_T g_\theta(\theta) = 0$ as seen above, we have by Taylor's formula, using the Hessian $Hg_\theta(T)$ of $g_\theta(T)$

$$\partial_T g_\theta(T) = \partial_T g_\theta(T) - \partial_T g_\theta(\theta)) \sim Hg_\theta(\theta)(T - \theta).$$

Since the MPLE estimator $\hat{\theta}_n$ verifies $\hat{G}(\hat{\theta}(n), n) = 0$, we have

$$Y(\hat{\theta}_n, n) = \sqrt{n}\ (\hat{G}(\hat{\theta}_n, n) - \partial_T g_\theta(\hat{\theta}_n)) = -\sqrt{n}\ \partial_T g_\theta(\hat{\theta}_n)).$$

The last two relations give then the approximation

$$-Y(\hat{\theta}_n, n) \sim \sqrt{n}\ Hg_\theta(\theta)\ (\hat{\theta}_n - \theta).$$

Since for n large, $\hat{\theta}_n$ is close to θ in probability, we see that with high probability $Y(\hat{\theta}_n, n)$ is close to $Y(\theta, n)$. We conclude that for n large, one has with high probability the approximation

$$-Y(\theta, n) \sim \sqrt{n}\ Hg_\theta(\theta)(\hat{\theta}_n - \theta).$$

Since the Hessian $\mathcal{H}(\theta) = Hg_\theta(\theta)$ is an invertible matrix, the random vector

$$\sqrt{n}\ (\hat{\theta}_n - \theta) \sim -\mathcal{H}(\theta)^{-1} Y(\theta, n)$$

must become, for n large, approximately Gaussian with mean zero and covariance matrix $\mathcal{H}(\theta)^{-1}\Sigma(\theta)\mathcal{H}(\theta)^{-1}$.

C. Estimated optimal separator

Fix two autologistic distributions π^+, π^- on $\mathcal{B}$ parameterized by θ^+, θ^- in Θ. Let $\hat{\theta}_n^+$ and $\hat{\theta}_n^-$ be the MPLE estimators of θ^+, θ^-, computed from two separate samples of n configurations resp. generated by π^+ and π^-. The optimal separator $f(x) =< u, U(x) > + a$ between π^+ and π^-, and its natural estimator $\hat{f}(x) =< \hat{u}, U(x) > + \hat{a}$ are given by formulas (7.10) and (7.11).

Proposition 3 In the preceding situation, as $n \to \infty$, the normalized error vector $\sqrt{n}(\hat{u} - u)$ is asymptotically Gaussian with mean zero and covariance matrix $Cov = \Gamma(\theta^+) + \Gamma(\theta^-)$, where the matrices $\Gamma(\theta)$ have been computed in Proposition 2. Moreover, $\sqrt{n}\,(\hat{a} - a)$ is also asymptotically Gaussian with mean zero and asymptotic variance $\mathrm{var}(\theta^+, \theta^-)$ verifying

$$\mathrm{var}(\theta^+, \theta^-) \leqslant \frac{\mathrm{L}(\mathrm{L}+1)}{2} \times \mathrm{trace}\left[\Gamma(\theta^+) + \Gamma(\theta^-)\right]. \tag{C.25}$$

Proof: For $\theta \in \Theta$, let $\varepsilon_n(\theta) = \sqrt{n}\,(\hat{\theta}_n - \theta)$ be the vector of normalized estimation errors. The independent vectors $\varepsilon_n(\theta^+)$ and $\varepsilon_n(\theta^-)$ are both asymptotically Gaussian with mean zero and covariance matrices $\Gamma(\theta^+)$, $\Gamma(\theta^-)$ computable via Proposition 2. Hence the vector $\sqrt{n}(\hat{u} - u)$ is asymptotically Gaussian with mean zero and covariance matrix $Cov = \Gamma(\theta^+) + \Gamma(\theta^-)$.

The partition function $Z(\theta) = \sum_{x \in \mathcal{B}} \exp(-\theta U(x))$ of π_θ is a smooth positive function of $\theta \in \Theta$. The vector valued function $\psi(\theta) = [\,\theta, z(\theta)\,]$ where $z(\theta) = \log(Z(\theta))$ is hence also a smooth in θ. As is well known (see for instance Azencott and Dacunha-Castelle (1985)), smooth functions preserve asymptotic normality. Hence since $\varepsilon_n(\theta)$ is asymptotically Gaussian with mean zero and covariance matrix $\Gamma(\theta)$, we conclude that $\sqrt{n}\,(\psi(\hat{\theta}_n) - \psi(\theta))$ must be asymptotically Gaussian with mean zero and covariance matrix $\psi'(\theta)^*\Gamma(\theta)\psi'(\theta)$, where the gradient $\psi'(\theta)$ is given by

$$\psi'(\theta)v = [\,v, v^* z'(\theta)\,] \quad \textit{for all vectors} \quad v \in \Theta.$$

In particular, $\sqrt{n}\,(z(\hat{\theta}_n) - z(\theta))$ is asymptotically Gaussian with mean zero and variance

$$\tau(\theta) = z'(\theta)^*\Gamma(\theta)z'(\theta),$$

where the column vector $z'(\theta)$ is the gradient of $z(\theta)$.

By a known trace identity, one has then the upper bound

$$\tau(\theta) \leqslant ||\, z'(\theta)\,||^2 \; \textit{trace}\,[\,\Gamma(\theta)\,].$$

The gradient of $Z(\theta)$ is given by

$$Z'(\theta) = -\sum_{x \in \mathcal{B}} U(x)\, \exp(- < \theta, U(x) >).$$

Since $U(x)$ has dimension $k = L(L+1)/2$, one has $||\, U(x)\,|| \leqslant \sqrt{k}$, and hence $||\, Z'(\theta)\,|| \leqslant \sqrt{k} Z(\theta)$.

This yields immediately $||\, z'(\theta)\,|| = ||\, Z'(\theta)/Z(\theta)\,|| \leqslant \sqrt{k}$ and thus finally

$$\tau(\theta) \leqslant k \times \textit{trace}\,[\,\Gamma(\theta)\,].$$

Applying this result to θ^+ and θ^- shows that $\sqrt{n}\,(\hat{a} - a)$ is asymptotically Gaussian with mean zero and variance verifying inequality (C.25). This concludes the proof.

Proposition 4 On the set $\mathcal{B}$ of binary vectors x, consider the two separators

$$f(x) = < u, U(x) > + a\,; \quad \phi(x) = < \eta, U(x) > + \alpha$$

parameterized by $u, \eta \in \Theta$ and $a, \alpha \in R$. Fix any autologistic distribution π_θ on $\mathcal{B}$. Let $p(f)$ and $p(\phi)$ be the two probabilities

$$p(f) = \pi_\theta[\, x \mid f(x) > 0\,] \quad \textit{and} \quad p(\phi) = \pi_\theta[\, x \mid \phi(x) > 0\,].$$

Recall that $k = L(L+1)/2$. Fix u, a, θ and assume that $0 < p(f) < 1$. For any arbitrary small percentage $0\% \leqslant \gamma < 100\%$, there is a strictly positive number q depending only on θ, u, a, γ such that for any pair (η, α) verifying

$$||\eta - u|| < \frac{q}{2\sqrt{k}} \quad \textit{and} \quad |\,\alpha - a\,| < q/2, \tag{C.26}$$

one must then have $|\,p(f) - p(\phi)\,| \leqslant \gamma$.

Proof: There exists a *strictly positive* number $q = q(\theta, u, a, \gamma)$ such that

$$\pi_\theta\,\{\, x \in \mathcal{B} \;\mid\; 0 < f(x) < q\,\} \;\leqslant\; \gamma\, p(f),$$
$$\pi_\theta\,\{\, x \in \mathcal{B} \;\mid\; -q < f(x) < 0\,\} \;\leqslant\; \gamma\,(1 - p(f)).$$

Indeed, since $\mathcal{B}$ is finite, this is already true for $\gamma = 0$, and hence a fortiori for any $\gamma \geqslant 0$. But for any given γ, we will systematically select q as the largest number verifying the preceding two inequalities.

Write $F = \{\, x \in \mathcal{B} \mid f(x) > 0\,\}$ and $\Phi = \{\, x \in \mathcal{B} \mid \phi(x) > 0\,\}$ so that $p(f) = \pi_\theta(F)$ and $p(\phi) = \pi_\theta(\Phi)$. By construction of q, one has then the inequalities

$$\pi_\theta(F - \Phi) \leqslant \gamma\, p(f) + \pi_\theta\;\{\, x \in \mathcal{B} \mid f(x) \geqslant q > 0 \;\; \textit{and} \;\; \phi(x) < 0\,\}, \tag{C.27}$$

$$\pi_\theta(\Phi - F) \leqslant \gamma\,(1 - p(f)) + \pi_\theta\;\{\, x \in \mathcal{B} \mid f(x) \leqslant -q < 0 \;\; \textit{and} \;\; \phi(x) > 0\,\}. \tag{C.28}$$

The two joint conditions $f(x) \geqslant q > 0$ and $\phi(x) < 0$ imply the elementary inequalities

$$q \leqslant f(x) - \phi(x) =< u - \eta,\, U(x) > + a - \alpha \leqslant\; \sqrt{k}||u - \eta|| + |\,a - \alpha\,|,$$

where $||.||$ is the norm in R^k and one uses the bound $||U(x)|| \leqslant \sqrt{k}$. Similarly the joint conditions $f(x) \leqslant -q < 0$ and $\phi(x) > 0$ imply

$$q \leqslant -f(x) + \phi(x) =< -u + \eta,$$
$$U(x) > +\alpha - a \;\leqslant\; \sqrt{k}||u - \eta|| + |\,(a - \alpha)\,|.$$

Hence for any pair (η, α) verifying the conditions (C.26), we see that the 2nd terms in both inequations (C.27) and (C.28) must be zero, which implies

$$\pi_\theta(F - \Phi) \leqslant \gamma\, p(f) \quad \textit{and} \quad \pi_\theta(\Phi - F)) \leqslant \gamma\,(1 - p(f)).$$

Since one clearly has

$$|\, p(f) - p(\phi)\,| \leqslant \pi_\theta(F - \Phi) + \pi_\theta(\Phi - F)),$$

we conclude that $|\, p(f) - p(\phi)\,| \leqslant \gamma$.

D. Control of discrimination errors due to separator estimation

Proposition 5 Fix $\theta^+, \theta^- \in \Theta$ and let $\pi^+ = \pi_{\theta^+}$ and $\pi^- = \pi_{\theta^-}$. Let $f(x) = uU(x) + a$ be the optimal separator between π^+ and π^-, given by equation (7.10). From two random samples of n configurations resp. generated by π^+ and π^-, one computes by equation (7.11) the estimator $\hat{f}(x) = \hat{u}U(x) + \hat{a}$ of $f(x)$. Let (p^+, p^-) and $(\hat{p}^+, \hat{p}^-)$ be the probabilities of correct decisions resp. achievable by the separators f and $\hat{f}$ for discrimination between π^+ and π^-. Assume that $0 < p^+ < 1$ and $0 < p^- < 1$.

Let γ and κ be two arbitrary small numbers verifying $0 \leqslant \gamma < 1$ and $0 < \kappa < 1$. Then one can find N such that for $n > N$ one has

$$P(\,|\,\hat{p}^+ - p^+\,| \leqslant \gamma \;\; \textit{and} \;\; |\,\hat{p}^- - p^-\,| \leqslant \gamma\,) \;\geqslant\; 1 - \kappa.$$

A practical estimate of N can be computed using formula (D.33) below.

Proof: Let $q > 0$ be the largest strictly positive number verifying

$$\pi^+\{\, x \in \mathcal{B} \;\mid\; 0 < f(x) < q\,\} \;\leqslant\; \gamma\, p^+,$$
$$\pi^+\{\, x \in \mathcal{B} \;\mid\; -q < f(x) < 0\,\} \;\leqslant\; \gamma\,(1 - p^+).$$

Fix a number $A > 0$, which will be explicitly selected below. For $n > 4kA^2/q^2$ we have the elementary inequalities

$$P(\,||\,\hat{u} - u\,|| \geqslant \frac{q}{2\sqrt{k}}\,) = P(\,\sqrt{n}\,||\,\hat{u} - u\,|| \geqslant \sqrt{n}\,\frac{q}{2\sqrt{k}}\,) \leqslant P(\,\sqrt{n}\,||\,\hat{u} - u\,|| \geqslant A\,).$$

Let $V \in \Theta$ be a random Gaussian vector with mean zero and covariance matrix Cov. The asymptotic normality results proved above now imply

$$\lim_{n\to\infty} P(\,||\,\hat{u}-u\,||\geqslant \frac{q}{2\sqrt{k}}\,)\leqslant \lim_{n\to\infty} P(\,\sqrt{n}\,||\,\hat{u}-u\,||\geqslant A\,)=P(\,||V||\geqslant A\,). \tag{D.29}$$

Let $\lambda_1 \geqslant \ldots \geqslant \lambda_k > 0$ be the eigenvalues of the covariance matrix Cov. By diagonalization of Cov, one finds a $(k\times k)$ orthogonal matrix M such that the coordinates of MV become $\big[\,\sqrt{\lambda_1}\,H_1,\ldots,\sqrt{\lambda_k}\,H_k\,\big]$, where the H_j are independent standard Gaussian random variables. Since $||MV||=||V||$, we can then write

$$P(\,||V||^2\geqslant A^2\,)=P(\,\sum_{j=1}^{k}\lambda_j H_j^2\geqslant A^2\,)\leqslant P(\,\lambda_1\sum_{j=1}^{k}\,H_j^2\geqslant A^2\,).$$

Fix any small $\kappa>0$. Since $\sum_{j=1}^{k}\,H_j^2$ has a χ^2 distribution with k degrees of freedom, let $Q(\kappa)$ be the $100\times(1-\kappa)\%$ percentile of the χ^2_k distribution. Now select and fix $A^2=Q(\kappa)\lambda_1$ in the last inequation to derive

$$P(\,||V||\geqslant A\,)\leqslant\kappa.$$

After imposing $n>2kA^2/q^2=4k\lambda_1\,Q(\kappa)/q^2$, inequation (D.29) now implies

$$\lim_{n\to\infty} P(\,||\,\hat{u}-u\,||\geqslant \frac{q}{2\sqrt{k}}\,)\leqslant\kappa. \tag{D.30}$$

Let $R(\kappa)$ be the $100\times(1-\kappa)\%$ percentile of the standard Gaussian distribution. After imposing $n>4R(\kappa)/q^2$, the asymptotic normality of $\sqrt{n}\,(\hat{a}-a)$ proved above immediately shows, that

$$\lim_{n\to\infty} P(\,|\,\hat{a}-a\,|\geqslant\frac{q}{2}\,)\leqslant\kappa. \tag{D.31}$$

Combining inequations (D.30) and (D.31) immediately yields

$$\lim_{n\to\infty} P(\,||\,\hat{u}-u\,||<\frac{q}{2\sqrt{k}}\quad \textit{and}\quad |\,\hat{a}-a\,|<\frac{q}{2}\,)\;\geqslant 1-2\kappa. \tag{D.32}$$

Let $\hat{p}^+=\pi^+\,[\,x\;\mid\;\hat{f}(x)>0\,]$ be the probability of good decisions under π^+ based on the estimated separator $\hat{f}(x)$. In view of inequation (D.32), we can now apply the Proposition 4 to conclude that

$$\lim_{n\to\infty} P(\,|\,\hat{p}^+-p^+\,|\leqslant\gamma\,)\;\geqslant 1-2\kappa.$$

In fact the asymptotic normality results used above are practically accurate in most concrete cases as soon as $n>50$. So we see that when the separator $\hat{f}$ is estimated from n observations

generated by π^+ and $n > N$ observations generated by π^-, with

$$N = \max \left\{ 50, n > 4R(\kappa)/q^2, 4k\lambda_1 \, Q(\kappa)/q^2 \right\}, \tag{D.33}$$

we can essentially assert that $P(\, | \hat{p}^+ - p^+ | \leqslant \gamma \,) \geqslant 1 - 2\kappa$.

Now to prove that

$$\lim_{n \to \infty} P(\, | \hat{p}^- - p^- | \leqslant \gamma \,) \geqslant 1 - 2\kappa,$$

one can define a new $q > 0$ as the largest number verifying the two equations

$$\pi^- \{ \, x \in \mathcal{B} \mid 0 < f(x) < q \, \} \leqslant \gamma \, p^-,$$

$$\pi^- \{ \, x \in \mathcal{B} \mid -q < f(x) < 0 \, \} \leqslant \gamma \, (1 - p^-),$$

and then apply the same arguments as above. This achieves the proof.

References

Alexe, G., Alexe, S., Liotta, L. A., Petricoin, E., Reiss, M. and Hammer, P. L. (2004). Ovarian cancer detection by logical analysis of proteomic data. *Proteomics* **4**, 766–783.

Assareh, A. and Moradi, M. H. (2007). Extracting efficient fuzzy if-then rules from mass spectra of blood samples to early diagnosis of ovarian cancer. *Proceedings of the 2007 IEEE Symposium on Computational Intelligence in Bioinformatics and Computational Biology*.

Azencott, R., Beri, A., Ghadyan, Y., Joseph, N., Lehalle, C. A. and Rowley, M. (2014). Realtime market microstructure analysis: online Transaction Cost Analysis . *Journal of Quantitative Finance* **14**, 1167–1185.

Azencott, R. and Dacunha-Castelle, D. (1985). *Time series of irregular observations*. New York: Springer.

Azencott, R. and Graffigne, C. (1992). Non-supervised segmentation using multi-level Markov random fields. *In Proceedings of the 11th IAPR Congress, Image Speech and Signal Analysis, La Haye*.

Ball, G., Mian, S., Holding, F., Allibone, R. O., Lowe, J., Ali, S., Li, G., McCardle, S., Ellis, I. O., Creaser, C. *and others*. (2002). An integrated approach utilizing artificial

neural networks and SELDI mass spectrometry for the classification of human tumours and rapid identification of potential biomarkers. *Bioinformatics* **18**, 395–404.

Besag, J. (1975). Statistical analysis of non-lattice data. *Statistician* **24**, 179–195.

Bigi, B. (2003). Using Kullback-Leibler distance for text categorization. *In Proceedings of the ECIR-2003*, 305–319. 2633 of Lecture Notes in Computer Science.

Bigi, B., Huang, Y. and Mori, R. D. (2003). Vocabulary and language model adaptation using information retrieval. *In Proceedings of the ECIR-2003*, 305–319. 2633 of Lecture Notes in Computer Science.

Brémaud, P. (1999). *Markov chains: Gibbs fields, Monte Carlo Simulation, and Queues*. New York: Springer-Verlag.

Chalmond, B. (2003). *Modeling and Inverse Problems in Image Analysis*. New York: Springer-Verlag.

Dagan, I., Lee, L. and Pereira, F. (1999). Similarity-based models of word cooccurrence probabilities. *Machine Learning* **34**, 43–69.

Eckel-Passow, J.E., Oberg, A. L. and Therneau, T. M. (2009). An insight into high-resolution mass-spectrometry data. *Biostatistics* **10**, 481–500.

Fan, R. E., Chang, K. W., Hsieh, C. J., Wang, X. R. and Lin, C. J. (2008). LIBLINEAR: A library for large linear classification. *J. Mach. Learn. Res.* **9**, 1871–1874. Software available at http://www.csie.ntu.edu.tw/ cjlin/liblinear.

Geisser, S. (1993). *Predictive inference*. New York: Chapman and Hall.

Hastie, T., Tibshirani, R. and Friedman, J. (2009). *The elements of statistical learning*. New York: Springer.

Izmirlian, G. (2004). Application of the random forest classification algorithm to a SELDI-TOF proteomics study in the setting of a cancer prevention trial. *Annals of the New York Academy of Sciences* **1020**, 154–174.

Jensen, J. L. and Kunsch, H. R. (1994). On asymptotic normality of pseudo likelihood estimates for pairwise interaction processes. *Annals of the Institute of Statistical Mathematics* **46**, 475–486.

Kong, A., Gupta, C., Ferrari, M., Agostini, M., Bedin, C., Bouamrani, A., Tasciotti, E. and Azencott, R. (2014). Biomarker signature discovery from mass spectrometry data. *IEEE/ACM Transactions on Computational Biology and Bioinformatics* **11**, 766–772.

Kullback, S. and Leibler, R. A. (1951). On information and sufficiency. *Annal of Mathematical Statistics* **22**, 79–86.

Li, L., Tang, H., Wu, Z., Gong, J., Gruidl, M., Zou, J., Tockman, M and Clark, R. A. (2004). Data mining techniques for cancer detection using serum proteomic profiling. *Artificial Intelligence in Medicine* **32**, 71–83.

Wang, Z. and Chang, Y. I. (2011). Marker selection via maximizing the partial area under the ROC curve of linear risk scores. *Biostatistics* **12**, 369–385.

Yasui, Y., Pepe, M., Thompson, M. L., Adam, B., Wright, G. L., Qu Jr, Y., Potter, J. D., Winget, M., Thornquist, M. and Feng, Z. (2003). A data-analytic strategy for protein biomarker discovery: profiling of high-dimensional proteomic data for cancer detection. *Biostatistics* **4**, 449–463.

Yu, J. S., Ongarello, S., Fiedler, R., Chen, X. W., Toffolo, G., Cobelli, C. and Trajanoski, Z. (2005). Ovarian cancer identification based on dimensionality reduction for high-throughput mass spectrometry data. *Bioinformatics* **21**, 2200–2209.

Zhu, W., Wang, X., Ma, Y., Rao, M., Glimm, J. and Kovach, J. S. (2003). Detection of cancer-specific markers amid massive mass spectral data. *Proc. Natl. Acad. Sci. USA* **100**, 14666–14671.